The Unreasonable Effectiveness of Deep Learning

Yann LeCun
Facebook AI Research &
Center for Data Science, NYU
yann@cs.nyu.edu
http://yann.lecun.com

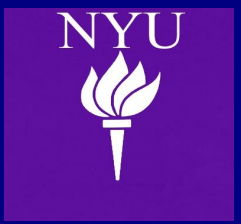

NYU Center for Data Science

A center devoted to the extraction of knowledge from data with the help of mathematics, statistics, and computation.

- Building on main campus (5th Avenue, 2016)
- We are looking for a new director!
- 9 open tenure-track faculty positions
- 5 “Data Science Fellow” positions
- Five-year grant from the Moore and Sloan Foundations
- Master of Science in Data Science program (started 2013)
- PhD in Data Science program (about to start)
- **http://cds.nyu.edu**

facebook

Facebook AI Research

At Facebook, we are advancing the field of machine intelligence and developing technologies that give people better ways to communicate. In the long term, we seek to understand intelligence and make intelligent machines.

Facebook AI Research (FAIR)

- A real research lab, open to the scientific community.
- 2 locations: Menlo Park and New York City
- 35 scientists and engineers (growing fast)
- 2 postdocs, many PhD interns
- Visiting faculty program
- Research projects spanning the spectrum from very long term to medium term to short term.
- Research themes: theory, applied math, optimization, data representation, unsupervised learning, learning sequential behavior, deep learning, structured prediction, NLP, speech, computer vision, software tools and infrastructure, parallel computing

http://research.facebook.com/ai

55 years of hand-crafted features

The traditional model of pattern recognition (since the late 50's)

- Fixed/engineered features (or fixed kernel) + trainable classifier

hand-crafted Feature Extractor → "Simple" Trainable Classifier →

Perceptron

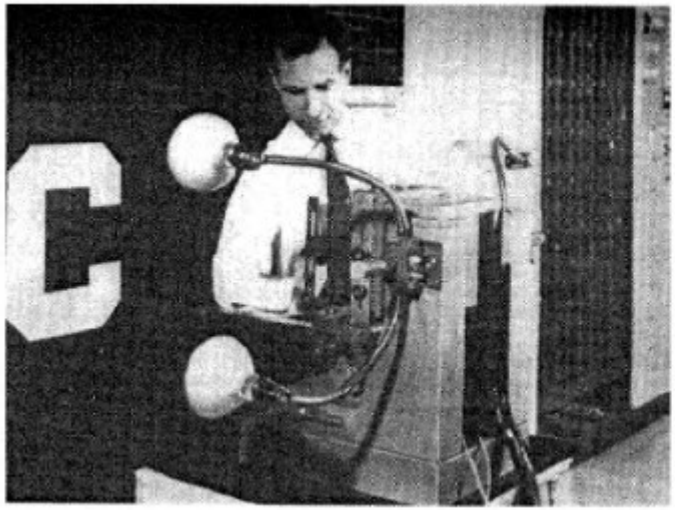

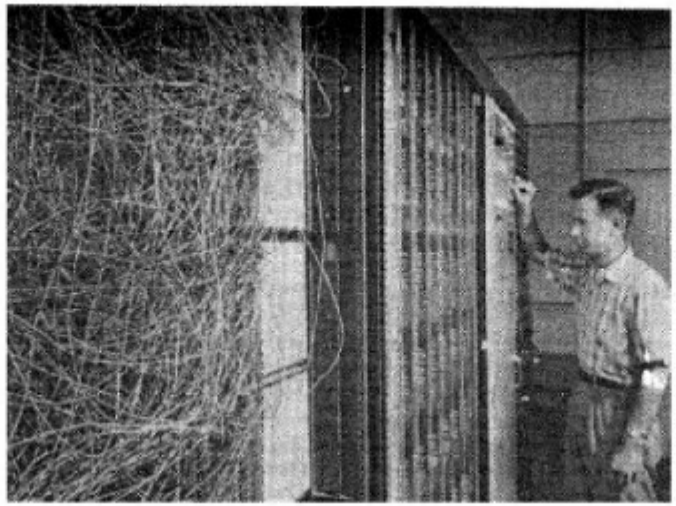

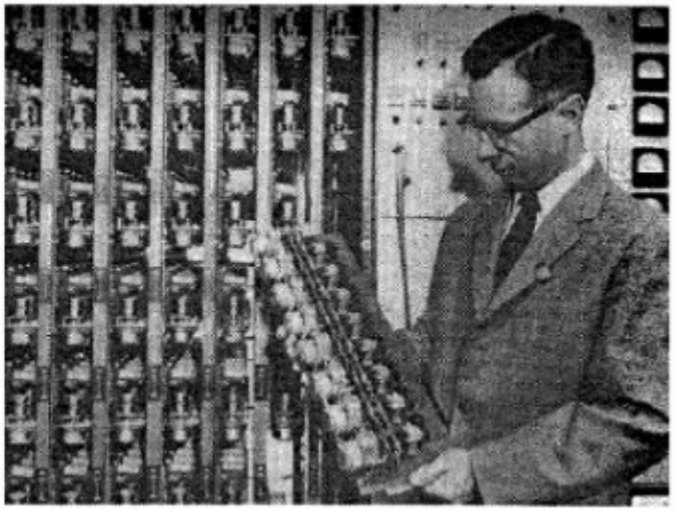

Architecture of "Classical" Recognition Systems

- "Classic" architecture for pattern recognition
  - Speech recognition: 1990-2011
  - Object Recognition: 2005-2012
  - Handwriting recognition (long ago)
  - Graphical model has latent variables (locations of parts)

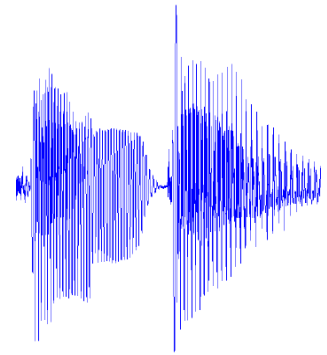

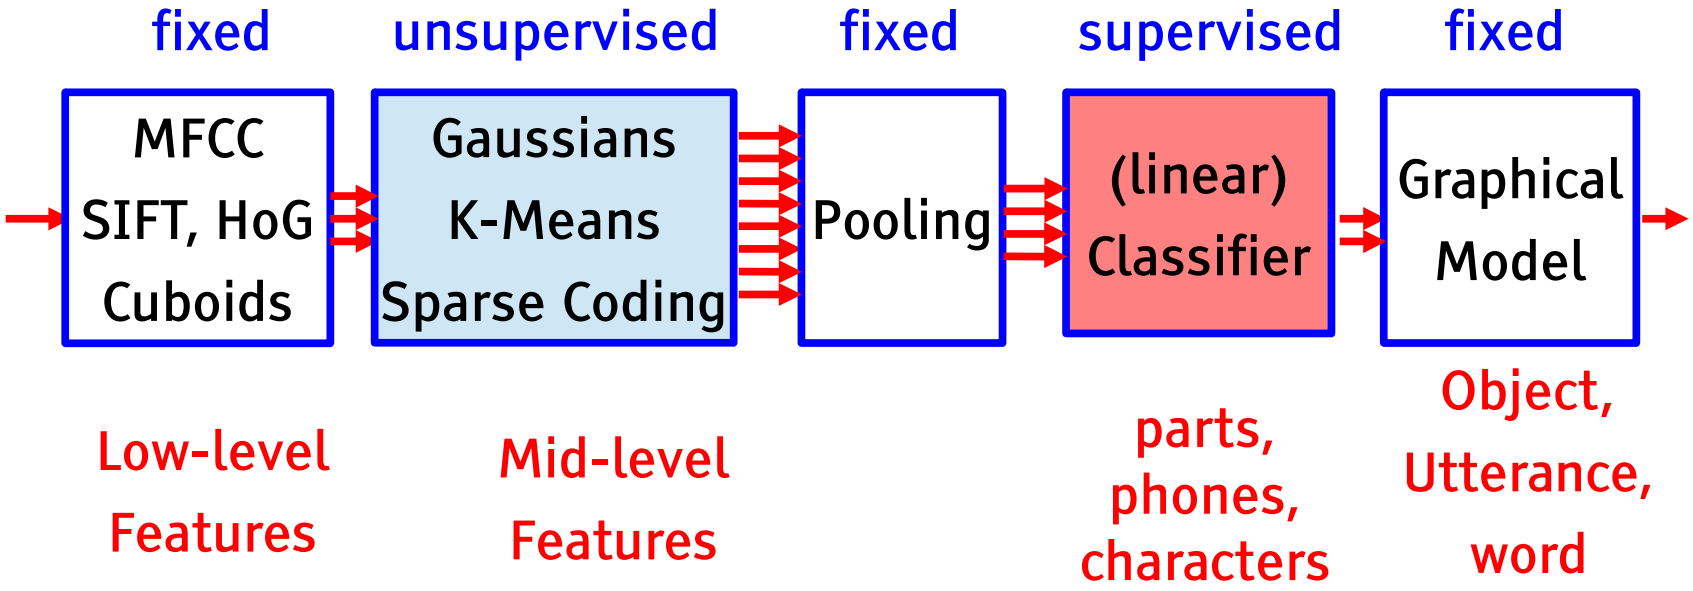

Architecture of Deep Learning-Based Recognition Systems

Y LeCun

- "Deep" architecture for pattern recognition
  - Speech, and Object recognition: since 2011/2012
  - Handwriting recognition: since the early 1990s
  - Convolutional Net with optional Graphical Model on top
  - Trained purely supervised
  - Graphical model has latent variables (locations of parts)

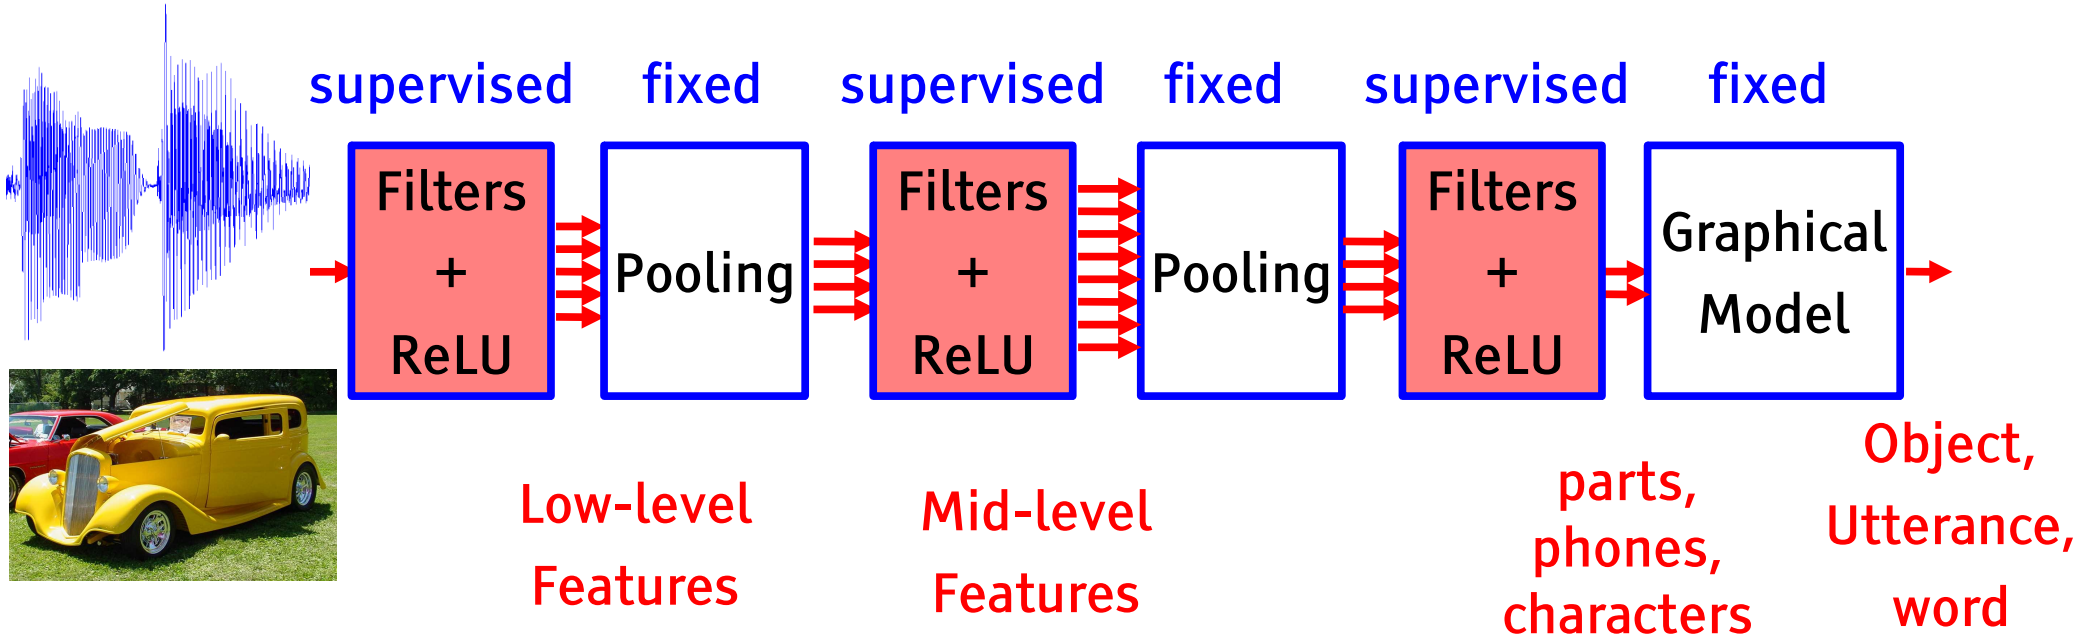

Future Systems: deep learning + structured prediction

- Globally-trained deep architecture
  - Handwriting recognition: since the mid 1990s
  - Speech Recognition: since 2011
  - All the modules are trained with a combination of unsupervised and supervised learning
  - **End-to-end training == deep structured prediction**

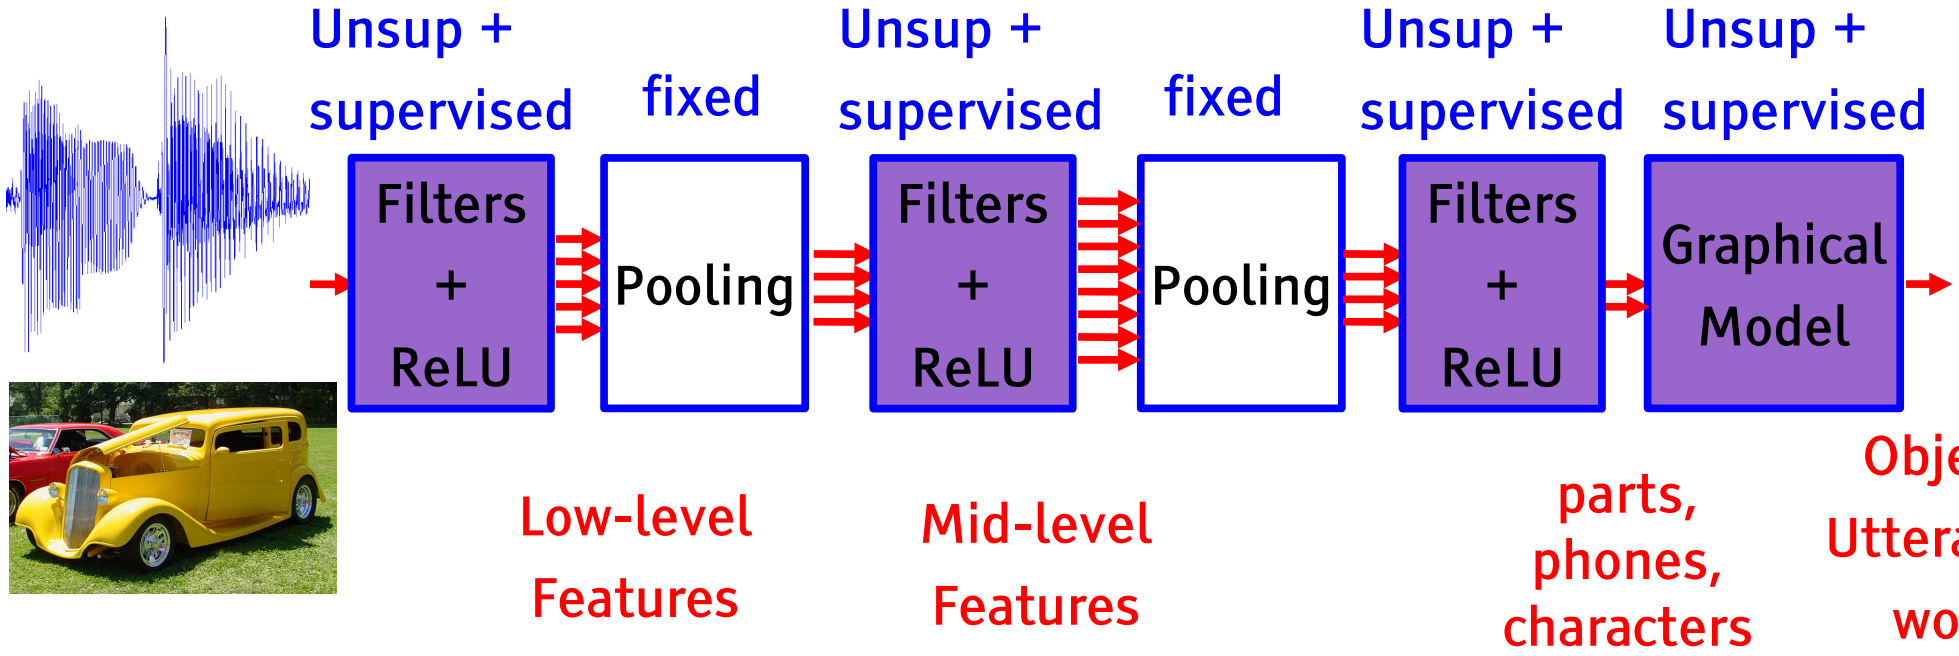

Deep Learning = Learning Hierarchical Representations

Y LeCun

- It's deep if it has more than one stage of non-linear feature transformation

Low-Level Feature → Mid-Level Feature → High-Level Feature → Trainable Classifier

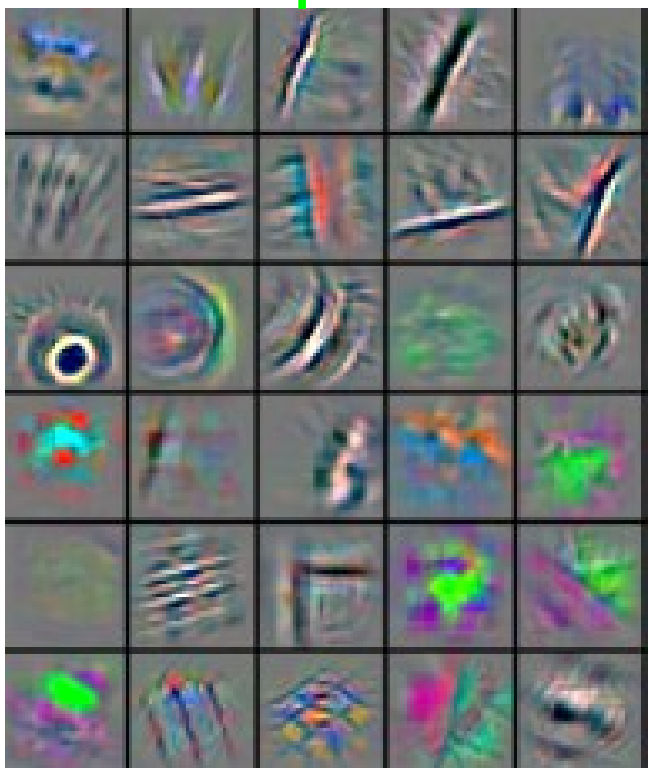

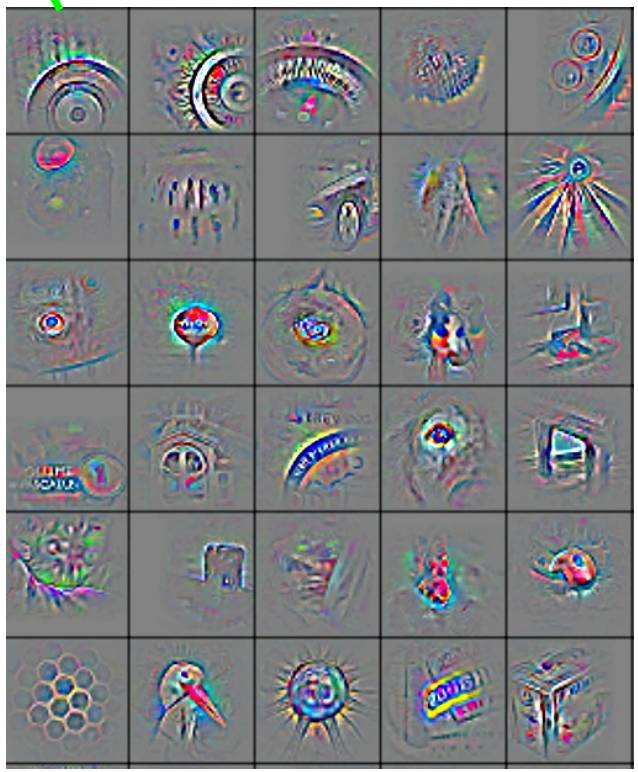

Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Trainable Feature Hierarchy

- Hierarchy of representations with increasing level of abstraction
- Each stage is a kind of trainable feature transform
- Image recognition
  - Pixel → edge → texton → motif → part → object
- Text
  - Character → word → word group → clause → sentence → story
- Speech
  - Sample → spectral band → sound → … → phone → phoneme → word

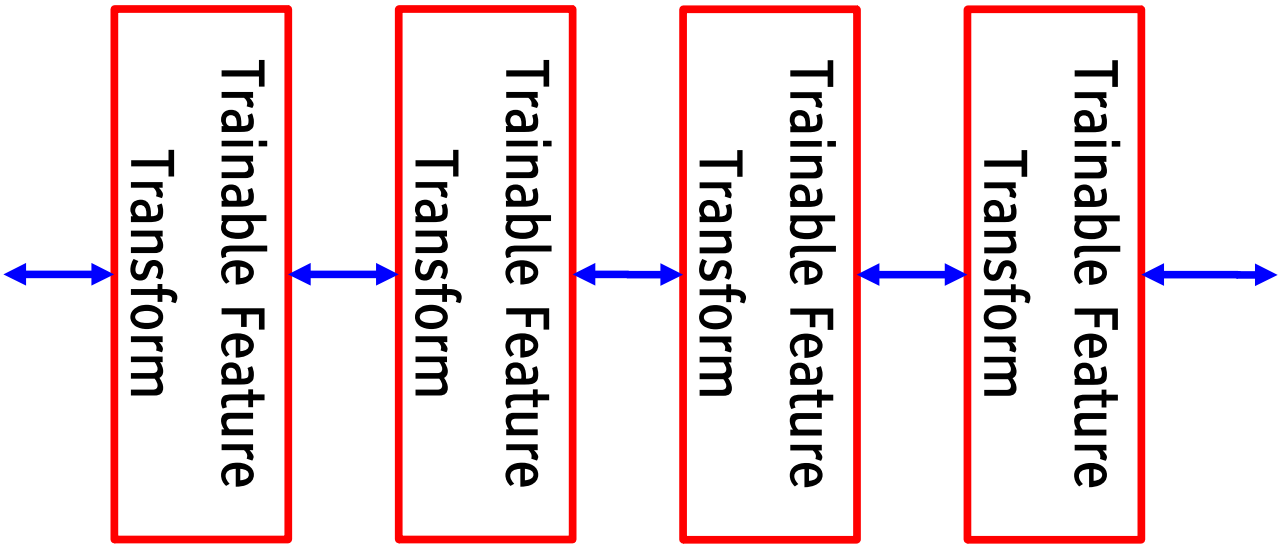

Learning Representations: a challenge for ML, CV, AI, Neuroscience, Cognitive Science...

Y LeCun

- How do we learn representations of the perceptual world?
  - How can a perceptual system build itself by looking at the world?
  - How much prior structure is necessary
- ML/AI: how do we learn features or feature hierarchies?
  - What is the fundamental principle? What is the learning algorithm? What is the architecture?
- Neuroscience: how does the cortex learn perception?
  - Does the cortex "run" a single, general learning algorithm? (or a small number of them)
- CogSci: how does the mind learn abstract concepts on top of less abstract ones?

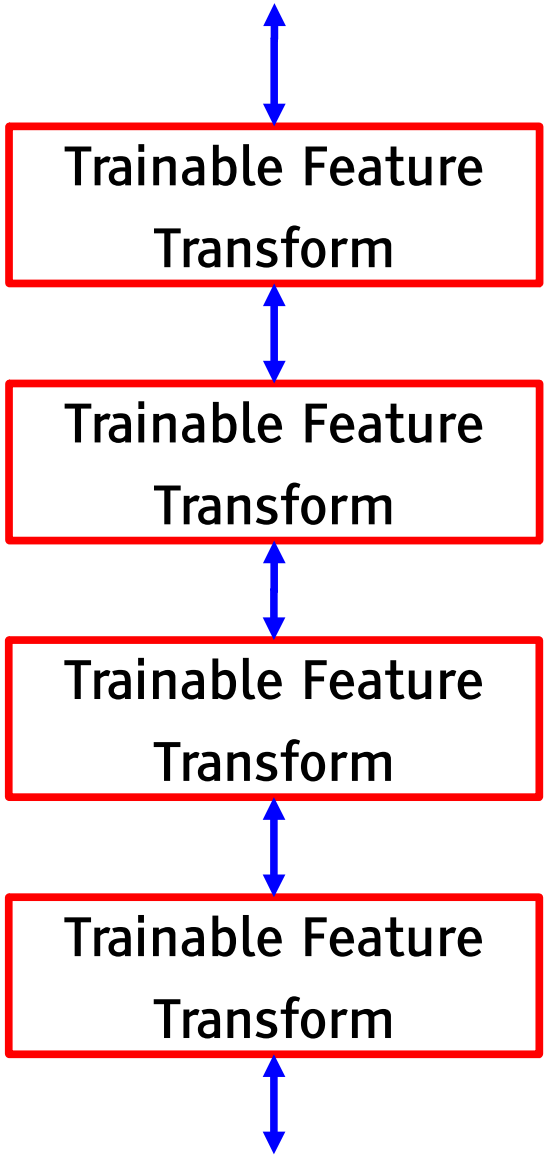

- Deep Learning addresses the problem of learning hierarchical representations with a single algorithm
  - or perhaps with a few algorithms

The Mammalian Visual Cortex is Hierarchical

- The ventral (recognition) pathway in the visual cortex has multiple stages
- Retina - LGN - V1 - V2 - V4 - PIT - AIT ....
- Lots of intermediate representations

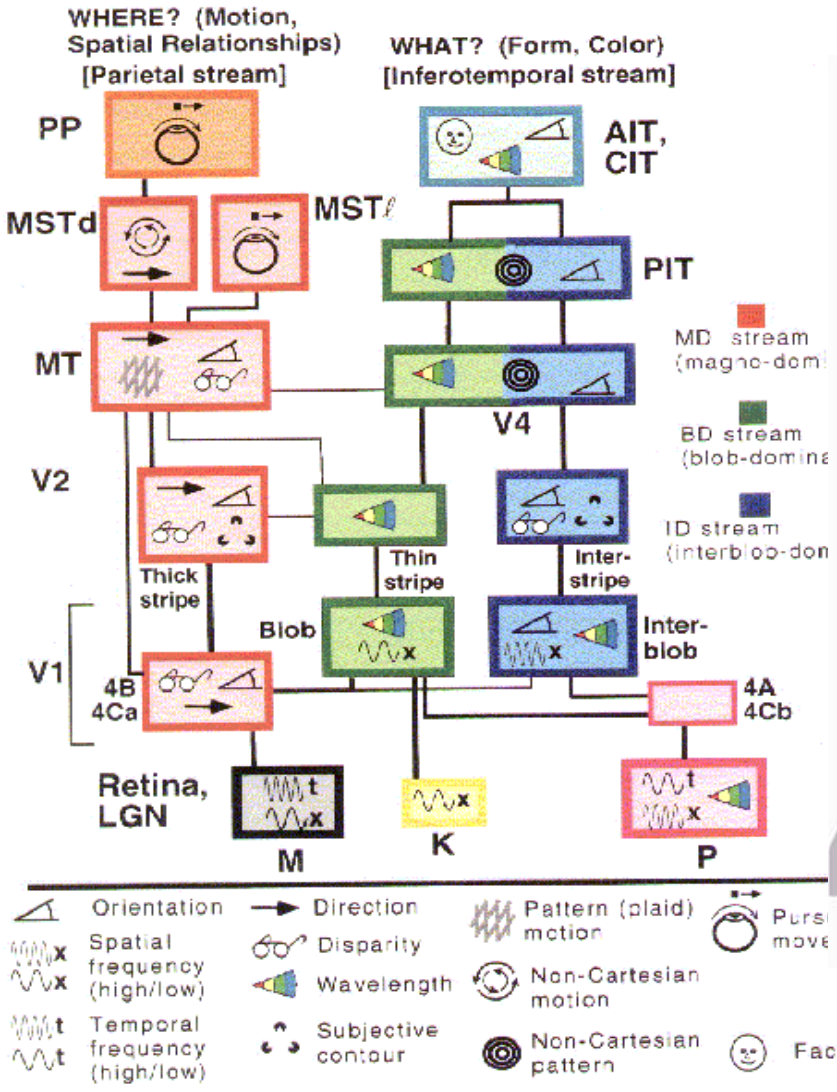

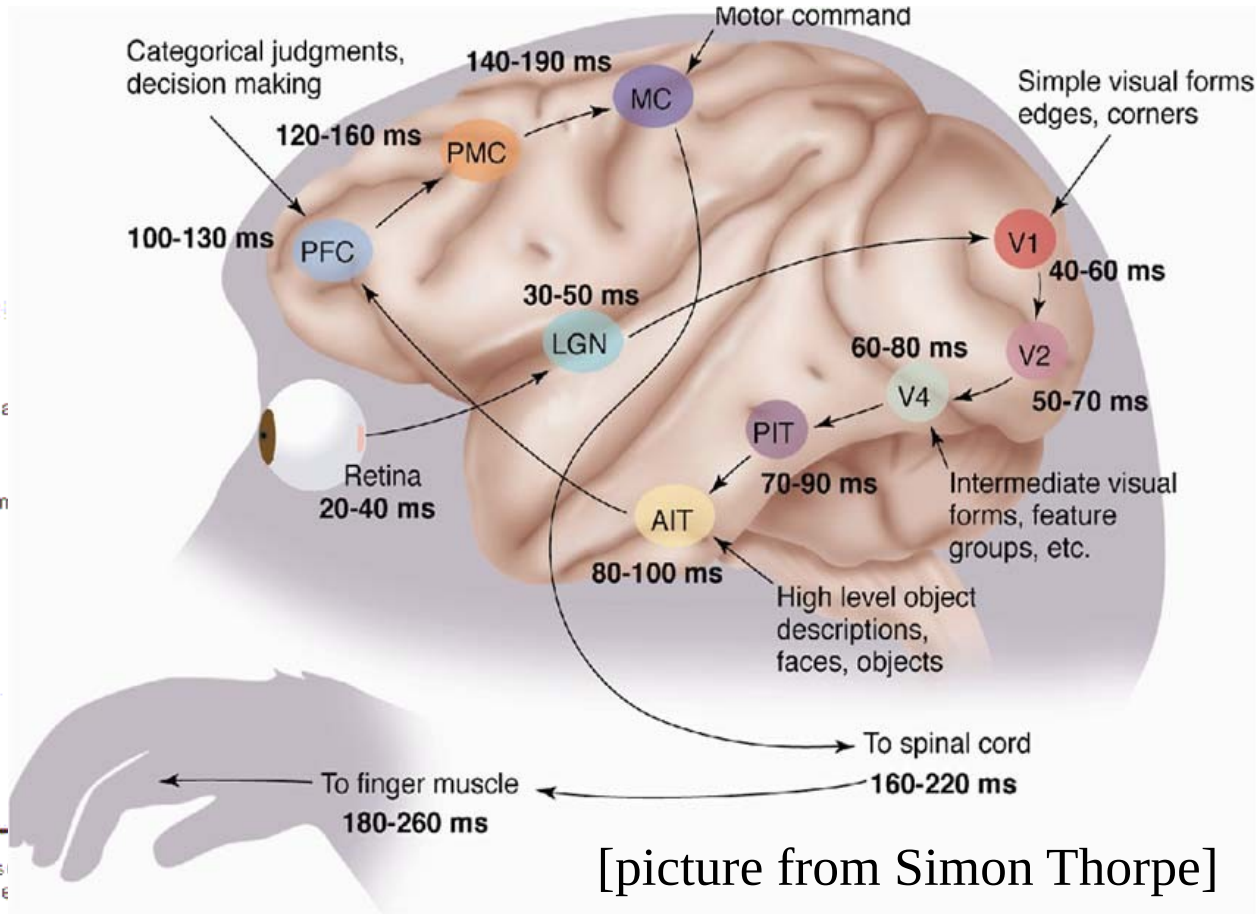

[picture from Simon Thorpe]

[Gallant & Van Essen]

Architecture of the Mammalian Visual Cortex

Y LeCun

- Ventral pathway = "what"
- dorsal pathway = "where"
- It's hierarchical
- There is feedback
- There is motion processing
- Learning is mostly unsupervised
- It does recognition, localization, navigation, grasping.....

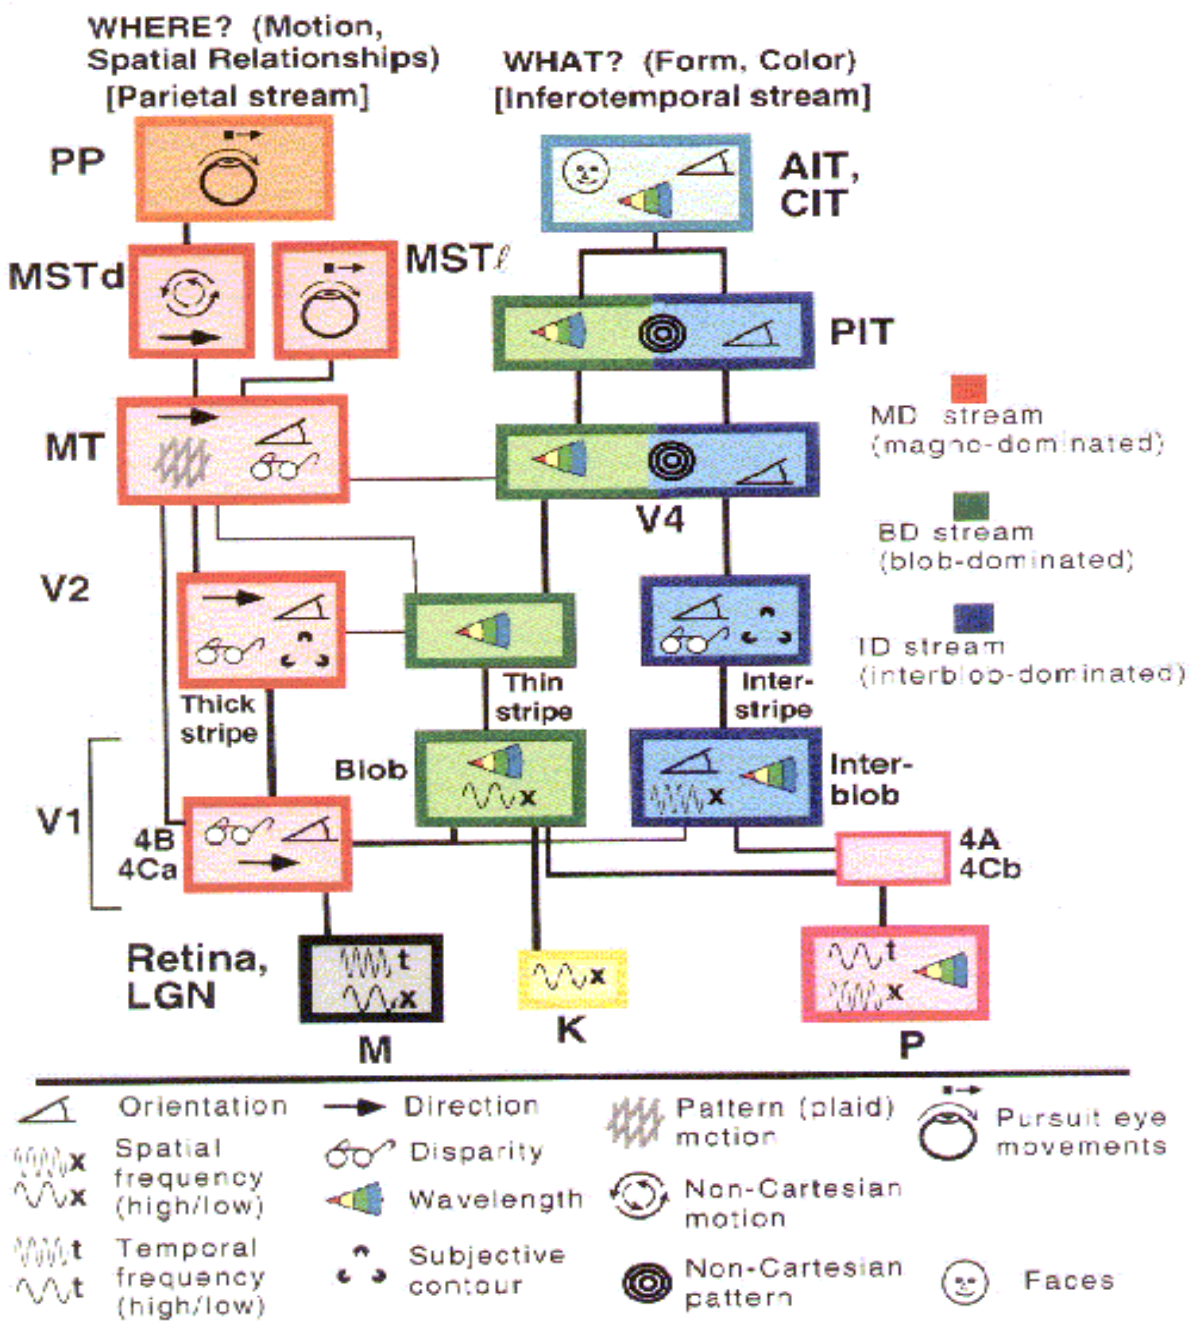

[Gallant & Van Essen]

Let's be inspired by nature, but not too much

Y LeCun

- It's nice imitate Nature,
- But we also need to understand
  - How do we know which details are important?
  - Which details are merely the result of evolution, and the constraints of biochemistry?
- For airplanes, we developed aerodynamics and compressible fluid dynamics.
  - We figured that feathers and wing flapping weren't crucial
- QUESTION: What is the equivalent of aerodynamics for understanding intelligence?

L'Avion III de Clément Ader, 1897
(Musée du CNAM, Paris)
His “Eole” took off from the ground in 1890,
13 years before the Wright Brothers, but you probably never heard of it (unless you are french).

Shallow vs Deep == lookup table vs multi-step algorithm

- "shallow & wide" vs "deep and narrow" == "more memory" vs "more time"
  - Look-up table vs algorithm
  - Few functions can be computed in two steps without an exponentially large lookup table
  - Using more than 2 steps can reduce the "memory" by an exponential factor.

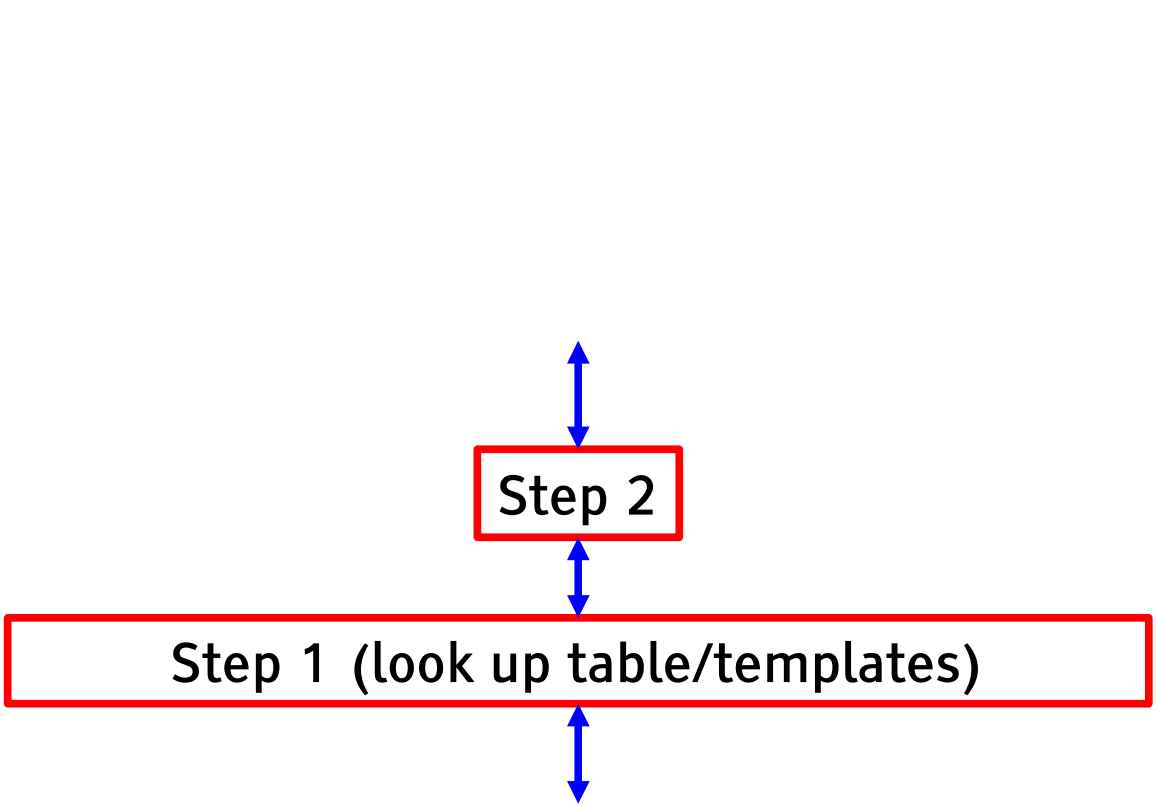

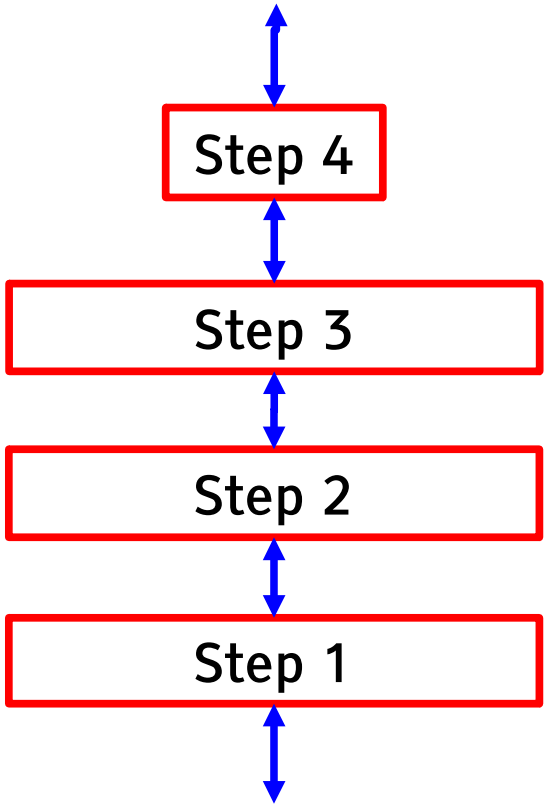

Which Models are Deep?

- 2-layer models are not deep (even if you train the first layer)
  - Because there is no feature hierarchy
- Neural nets with 1 hidden layer are not deep
- SVMs and Kernel methods are not deep
  - Layer1: kernels; layer2: linear
  - The first layer is “trained” in with the simplest unsupervised method ever devised: using the samples as templates for the kernel functions.
  - “glorified template matching”
- Classification trees are not deep
  - No hierarchy of features. All decisions are made in the input space

$$G(X,\alpha) = \sum_j \alpha_j K(X^j, X)$$

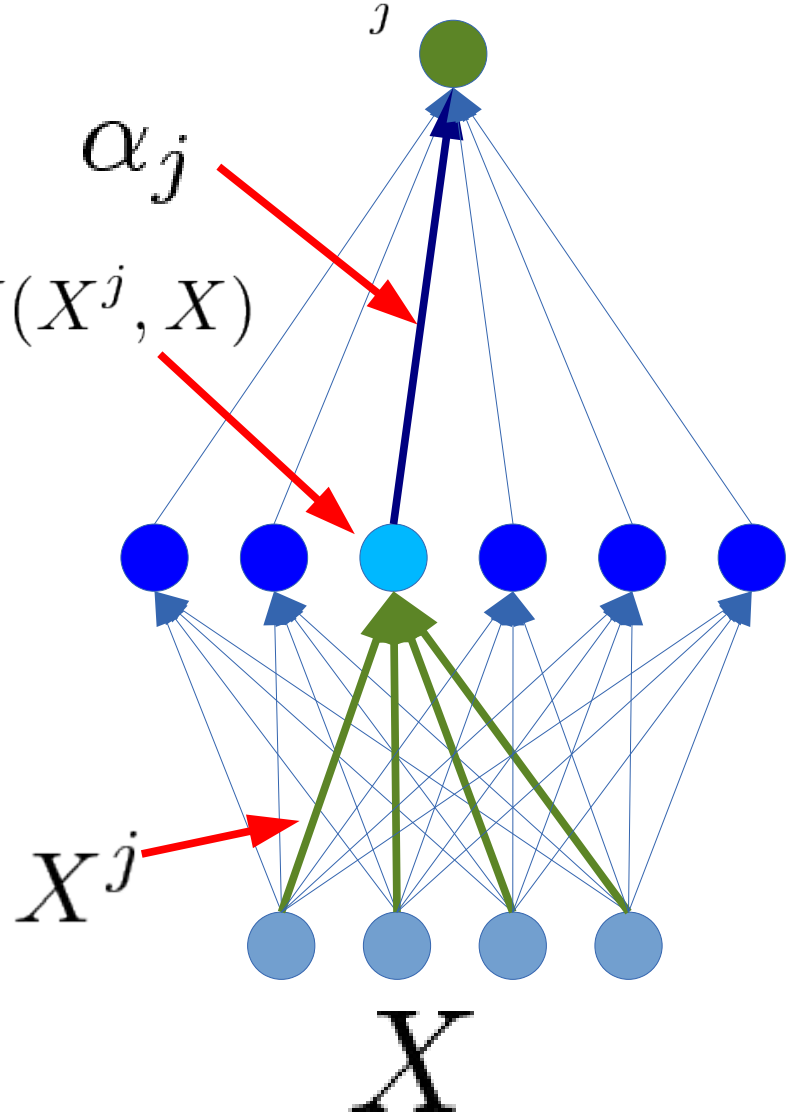

What Are
Good Feature?

Discovering the Hidden Structure in High-Dimensional Data
The manifold hypothesis

- Learning Representations of Data:
  - **Discovering & disentangling the independent explanatory factors**
- The Manifold Hypothesis:
  - Natural data lives in a low-dimensional (non-linear) manifold
  - Because variables in natural data

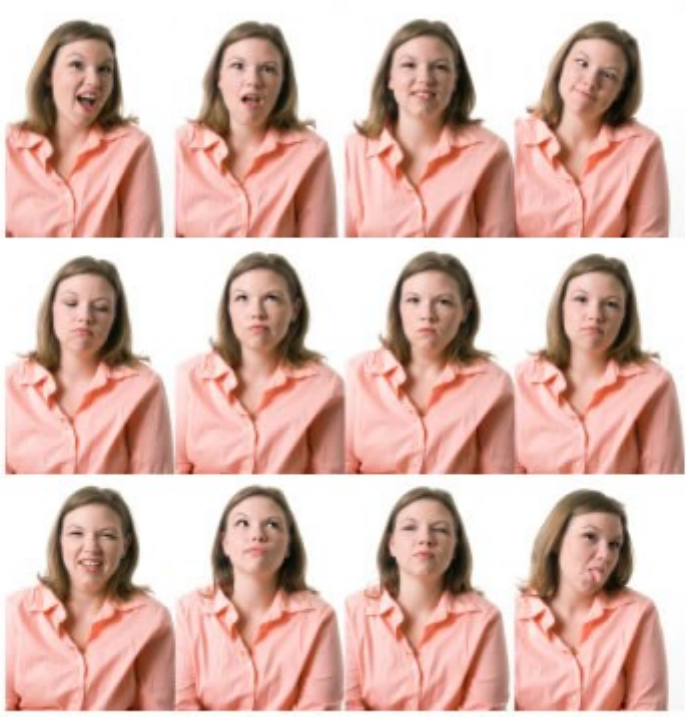

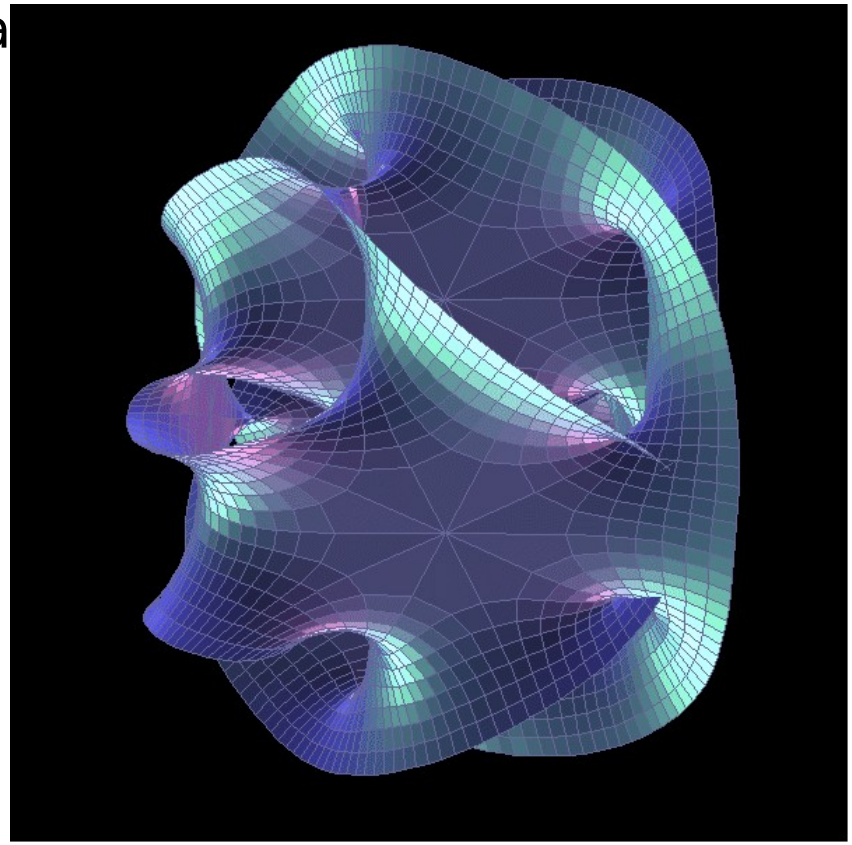

Basic Idea for Invariant Feature Learning

- **Embed the input non-linearly into a high(er) dimensional space**
  - In the new space, things that were non separable may become separable
- **Pool regions of the new space together**
  - Bringing together things that are semantically similar. Like pooling.

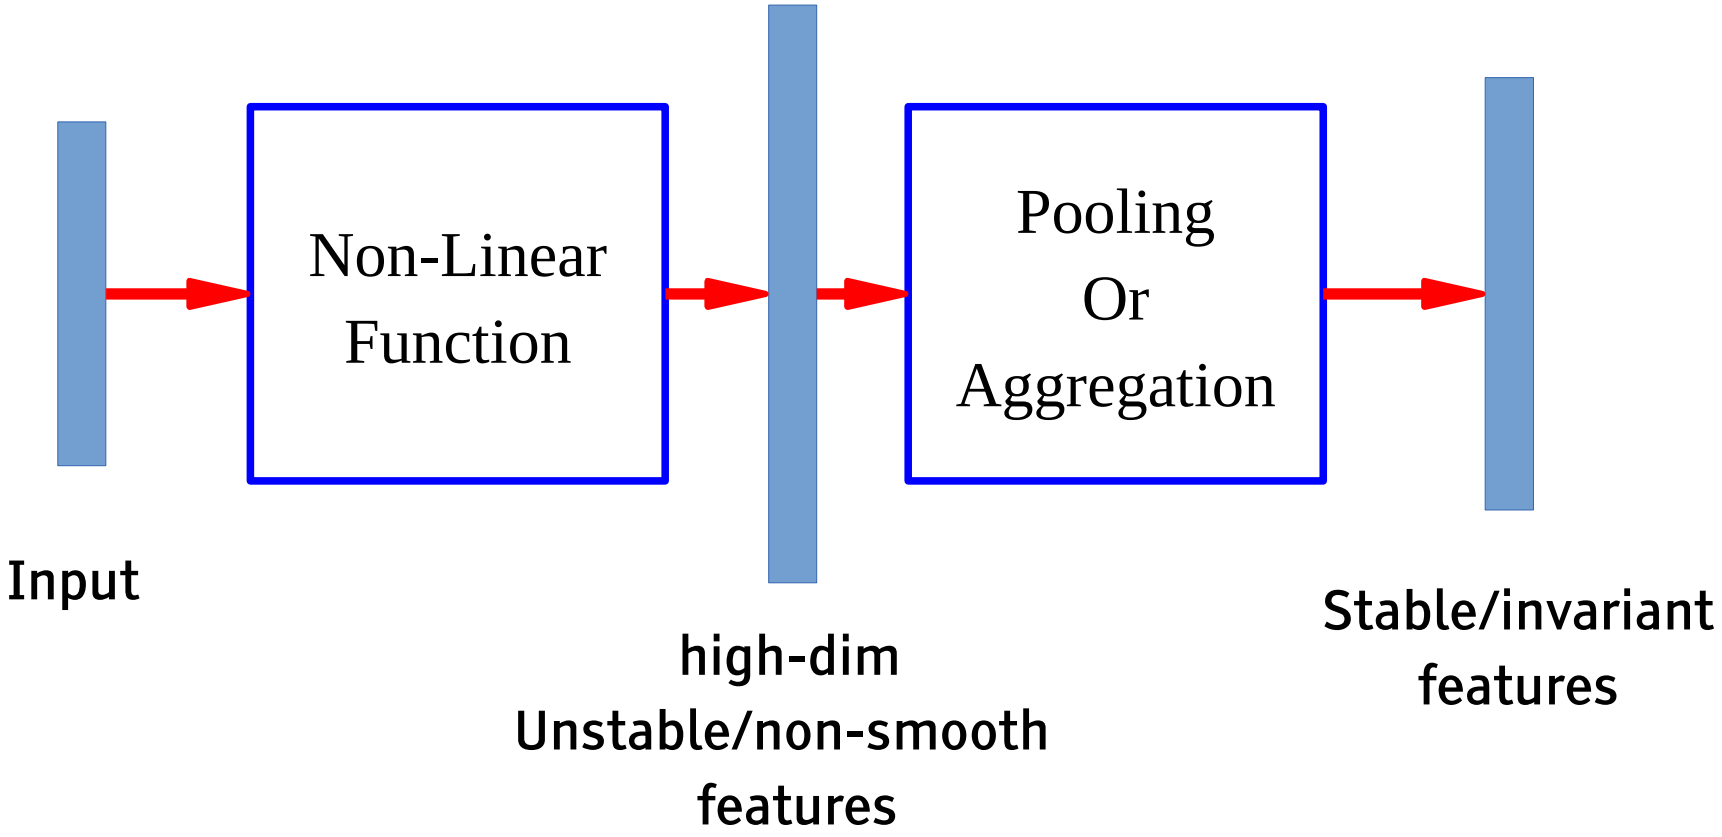

Sparse Non-Linear Expansion → Pooling

- Use clustering to break things apart, pool together similar things

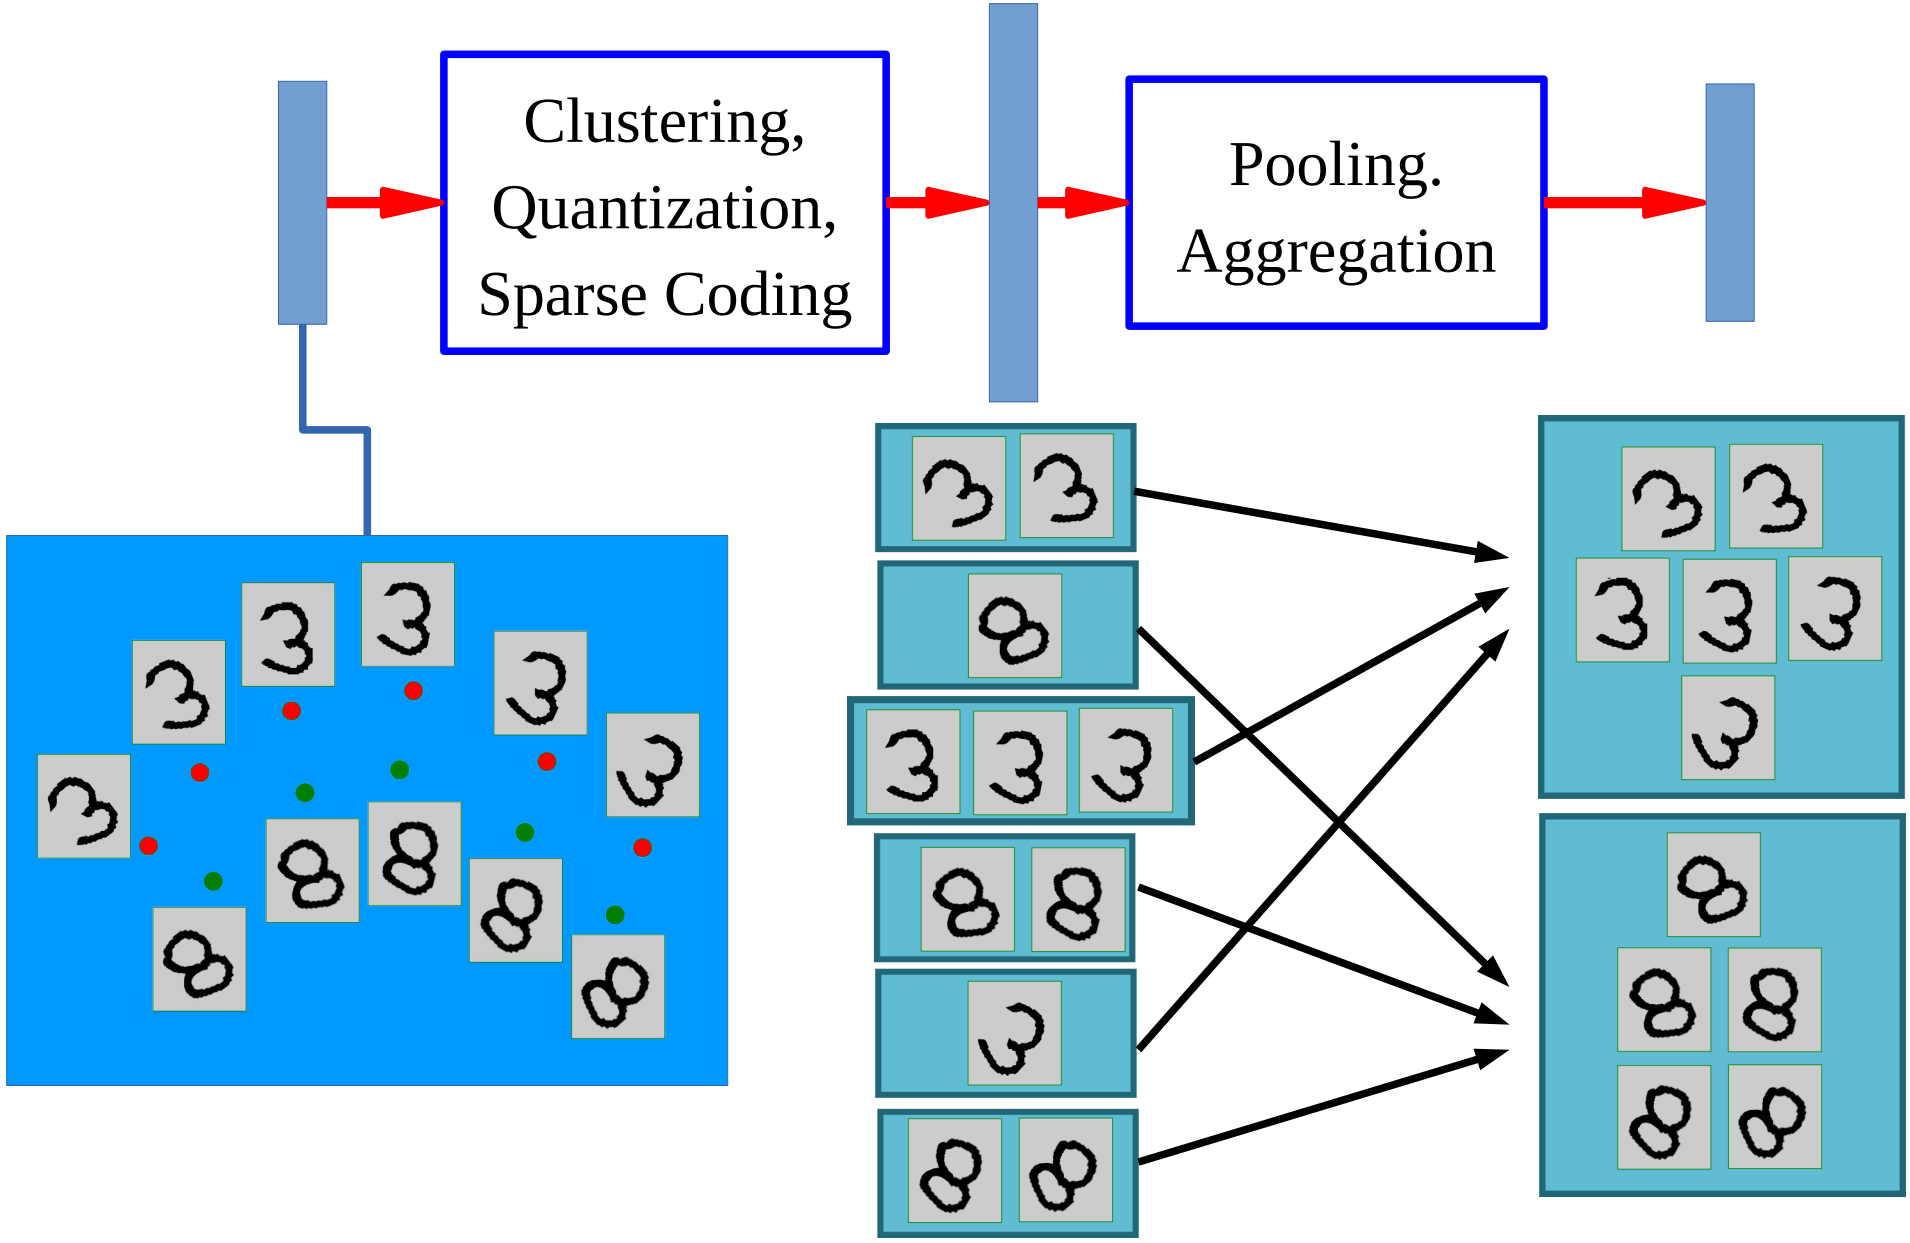

Overall Architecture: multiple stages of Normalization → Filter Bank → Non-Linearity → Pooling

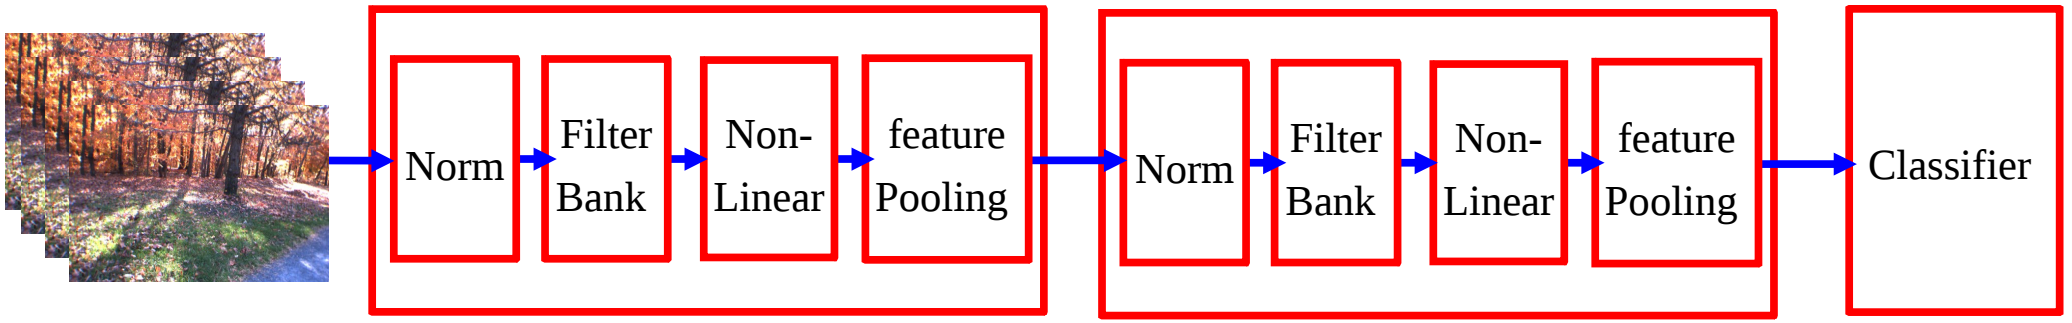

- **Normalization:** variation on whitening (optional)
  - Subtractive: average removal, high pass filtering
  - Divisive: local contrast normalization, variance normalization
- **Filter Bank:** dimension expansion, projection on overcomplete basis
- **Non-Linearity:** sparsification, saturation, lateral inhibition....
  - Rectification (ReLU), Component-wise shrinkage, tanh,..

$$ReLU(x) = max(x, 0)$$

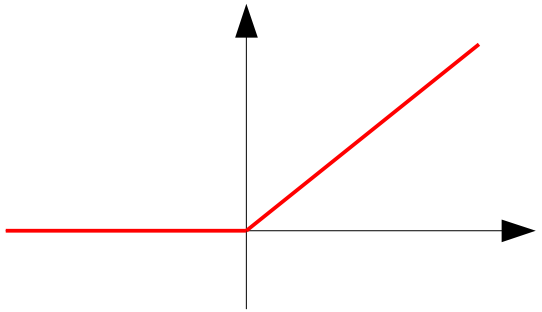

- **Pooling:** aggregation over space or feature type
  - Max, Lp norm, log prob.

$$MAX: Max_i(X_i); \quad L_p: \sqrt[p]{X_i^p}; \quad PROB: \frac{1}{b}\log\left(\sum_i e^{bX_i}\right)$$

Deep Nets with ReLUs and Max Pooling

- Stack of linear transforms interspersed with Max operators
- Point-wise ReLUs:

$$ReLU(x) = max(x, 0)$$

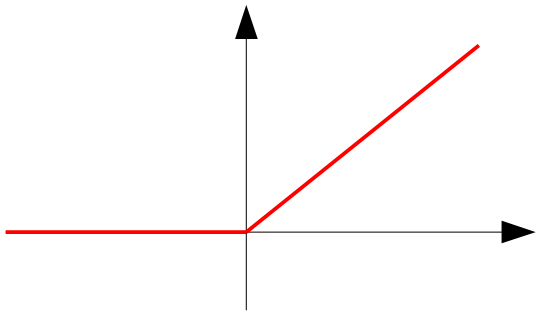

- Max Pooling
  - "switches" from one layer to the next

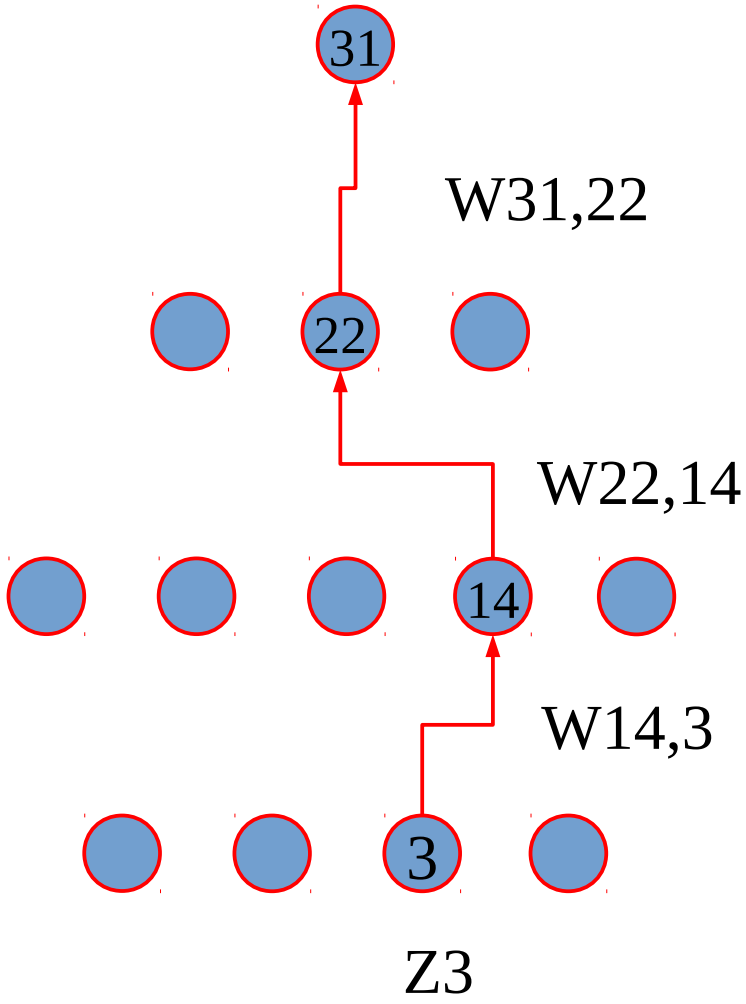

Supervised Training: Stochastic (Sub)Gradient Optimization

To compute all the derivatives, we use a backward sweep called the **back-propagation algorithm** that uses the recurrence equation for $\frac{\partial E}{\partial X_i}$

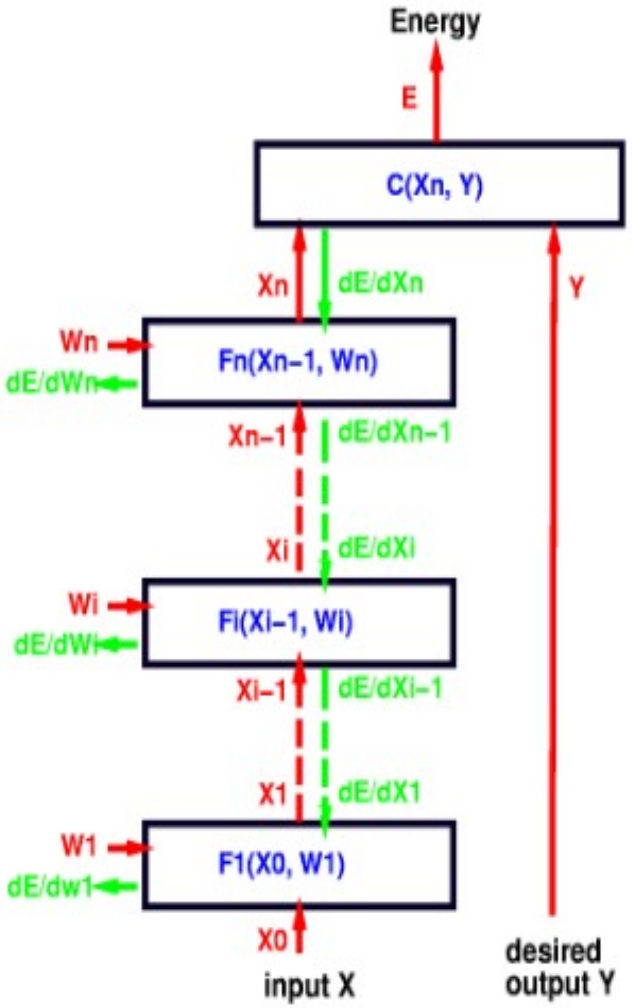

- $\frac{\partial E}{\partial X_n} = \frac{\partial C(X_n, Y)}{\partial X_n}$
- $\frac{\partial E}{\partial X_{n-1}} = \frac{\partial E}{\partial X_n} \frac{\partial F_n(X_{n-1}, W_n)}{\partial X_{n-1}}$
- $\frac{\partial E}{\partial W_n} = \frac{\partial E}{\partial X_n} \frac{\partial F_n(X_{n-1}, W_n)}{\partial W_n}$
- $\frac{\partial E}{\partial X_{n-2}} = \frac{\partial E}{\partial X_{n-1}} \frac{\partial F_{n-1}(X_{n-2}, W_{n-1})}{\partial X_{n-2}}$
- $\frac{\partial E}{\partial W_{n-1}} = \frac{\partial E}{\partial X_{n-1}} \frac{\partial F_{n-1}(X_{n-2}, W_{n-1})}{\partial W_{n-1}}$
- ....etc, until we reach the first module.
- we now have all the $\frac{\partial E}{\partial W_i}$ for $i \in [1, n]$.

Loss Function for a simple network

1-1-1 network

- Y = W1*W2*X

trained to compute the identity function with quadratic loss

- Single sample X=1, Y=1 L(W) = (1-W1*W2)^2

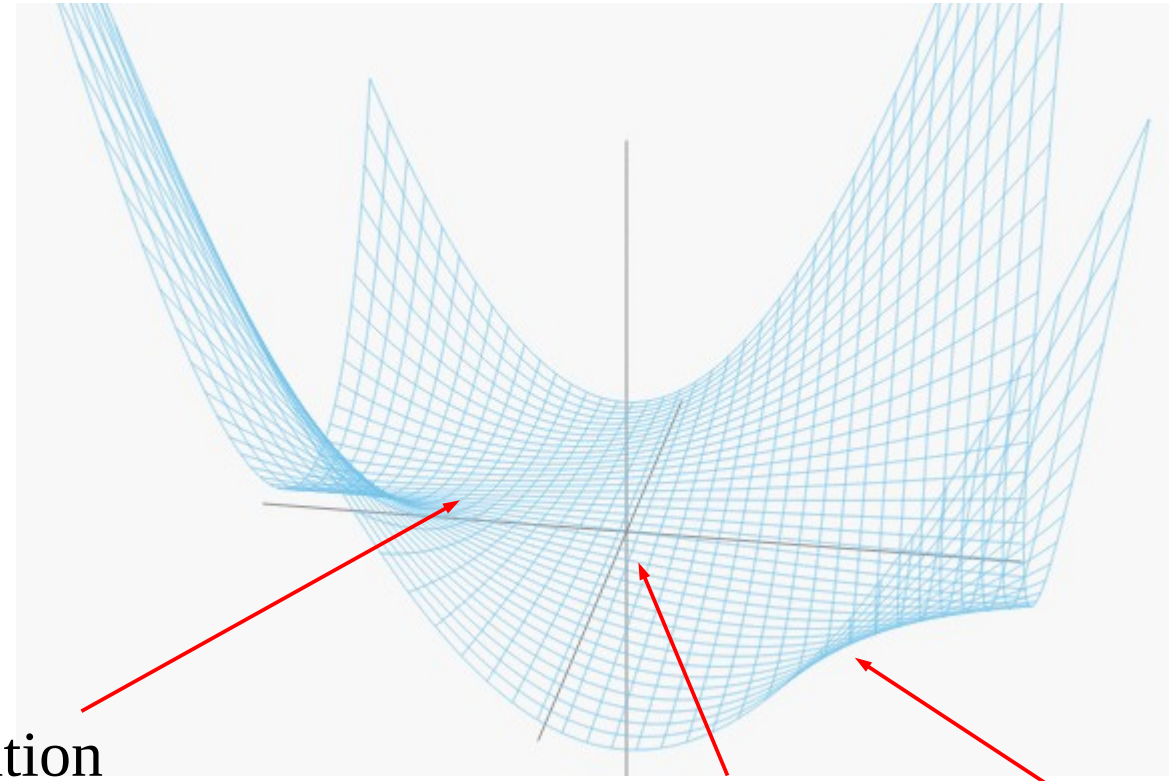

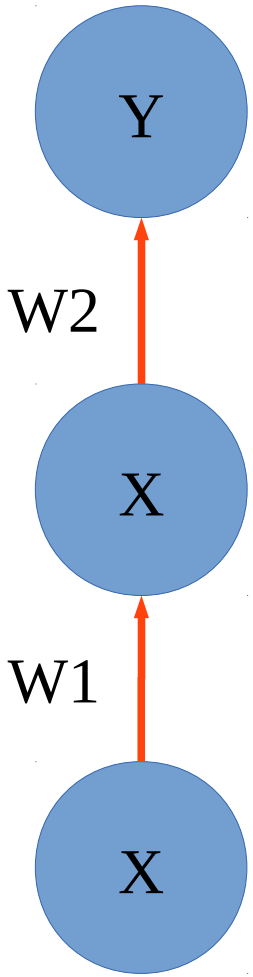

Deep Nets with ReLUs

- Single output:

$$\hat{Y} = \sum_P \delta_P(W, X)\left(\prod_{(ij)\in P} W_{ij}\right) X_{P_{start}}$$

- Wij: weight from j to i
- P: path in network from input to output
  - P=(3,(14,3),(22,14),(31,22))
- di: 1 if ReLU i is linear, 0 if saturated.
- Xpstart: input unit for path P.

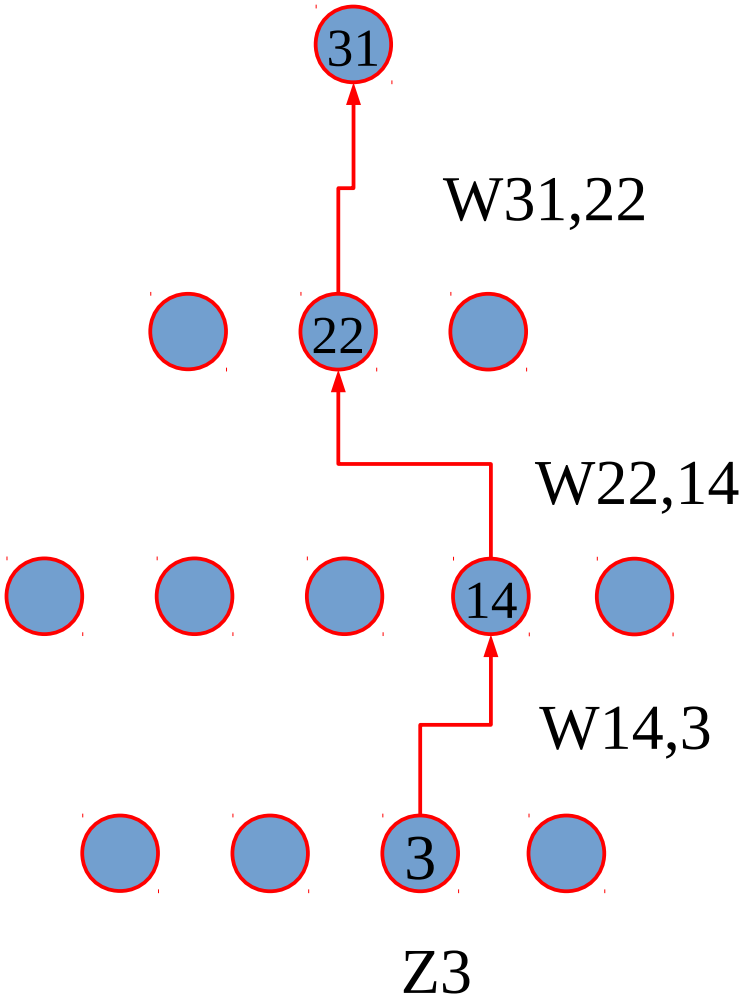

$$\hat{Y} = \sum_P \delta_P(W, X)\left(\prod_{(ij)\in P} W_{ij}\right) X_{P_{start}}$$

- Dp(W,X): 1 if path P is "active", 0 if inactive
- Input-output function is piece-wise linear
- Polynomial in W with random coefficients

Deep Convolutional Nets (and other deep neural nets)

- Training sample: (Xi,Yi) k=1 to K
- Objective function (with margin-type loss = ReLU)

$$L(W)=\sum_k ReLU\left(1-Y^k\sum_P \delta_P(W,X^k)\left(\prod_{(ij)\in P} W_{ij}\right)X^k_{P_{start}}\right)$$

$$L(W)=\sum_k\sum_P \left(X^k_{P_{start}}Y^k\right)\delta_P(W,X^k)\left(\prod_{(ij)\in P} W_{ij}\right)$$

$$L(W)=\sum_P \left[\sum_k \left(X^k_{P_{start}}Y^k\right)\delta_P(W,X^k)\right]\left(\prod_{(ij)\in P} W_{ij}\right)$$

$$L(W)=\sum_P C_p(X,Y,W)\left(\prod_{(ij)\in P} W_{ij}\right)$$

- Polynomial in W of degree l (number of adaptive layers)
- Continuous, piece-wise polynomial with "switched" and partially random coefficients
  - Coefficients are switched in an out depending on W

Deep Nets with ReLUs: Objective Function is Piecewise Polynomial

- If we use a hinge loss, delta now depends on label Yk:

$$L(W)=\sum_{P} C_p(X,Y,W)\left(\prod_{(ij)\in P} W_{ij}\right)$$

- Piecewise polynomial in W with random coefficients
- A lot is known about the distribution of critical points of polynomials on the sphere with random (Gaussian) coefficients [Ben Arous et al.]
  - High-order spherical spin glasses
  - Random matrix theory

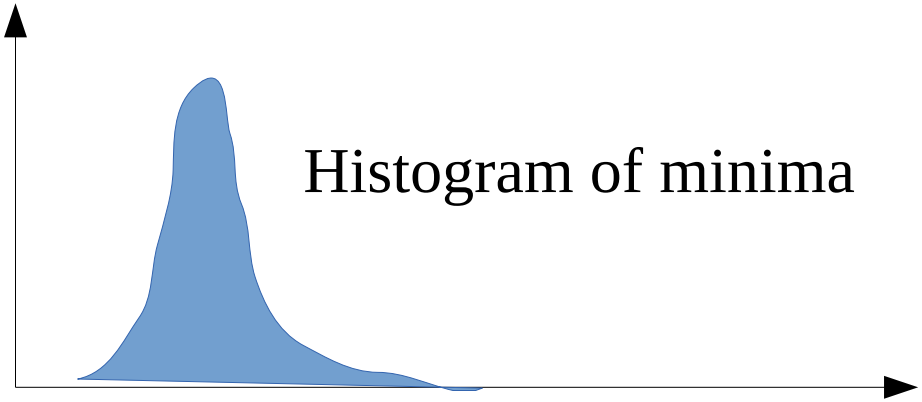

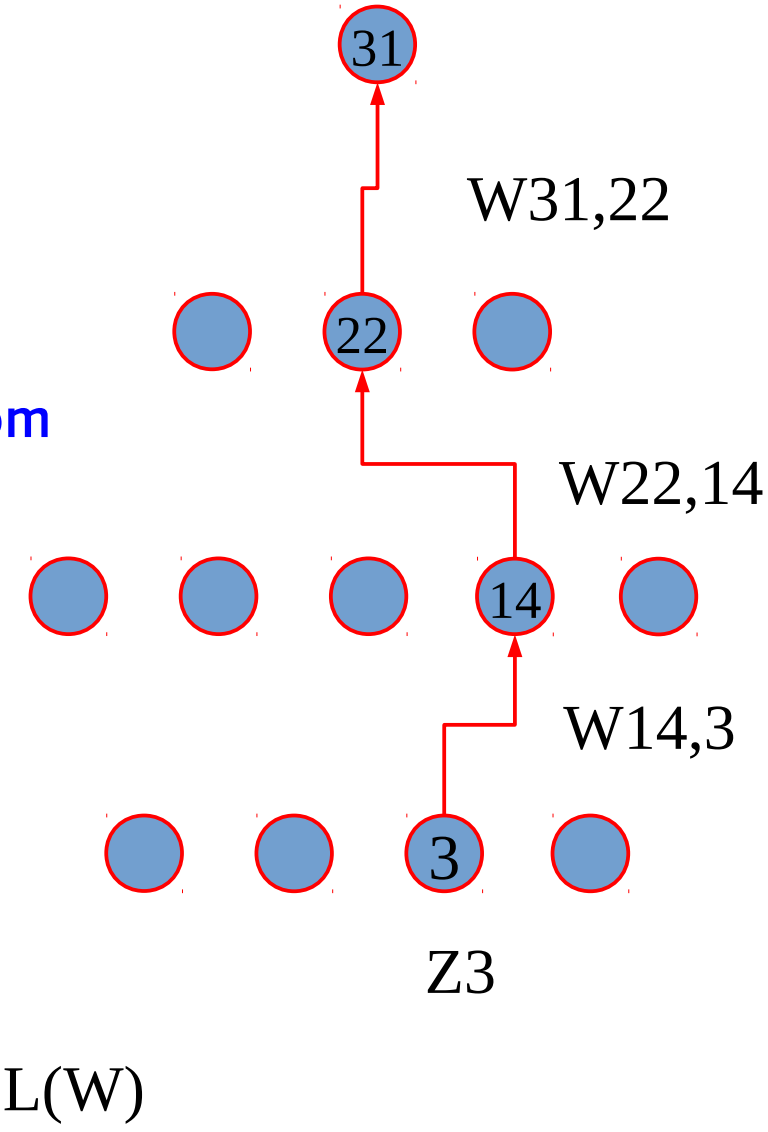

Deep Nets with ReLUs: Objective Function is Piecewise Polynomial

- Train 2-layer nets on scaled-down MNIST (10x10) from multiple initial conditions. Measure loss on test set.

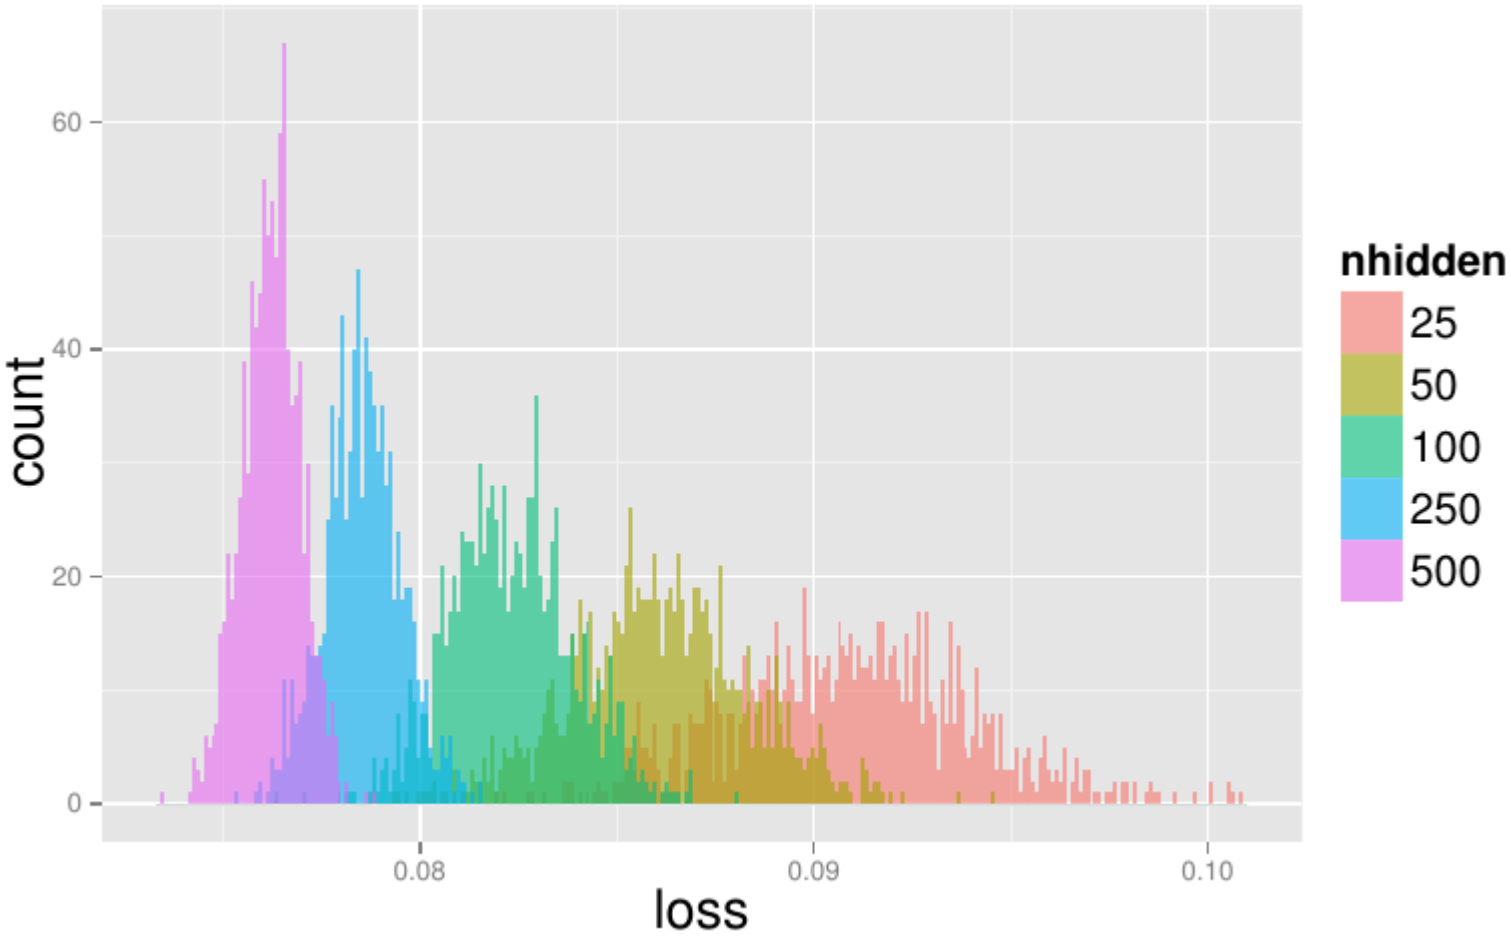

[Choromanska, Henaff, Mathieu, Ben Arous, LeCun 2015]

Spherical Spin Glass theory

Distribution of critical points (saddle points, minima, maxima)

- K=number of negative eigenvalues of Hessian (K=0 → minimum)

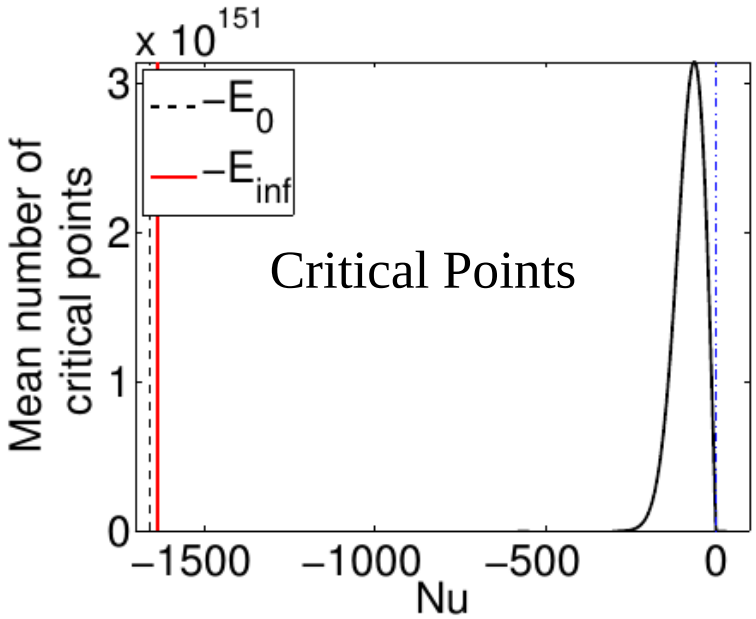

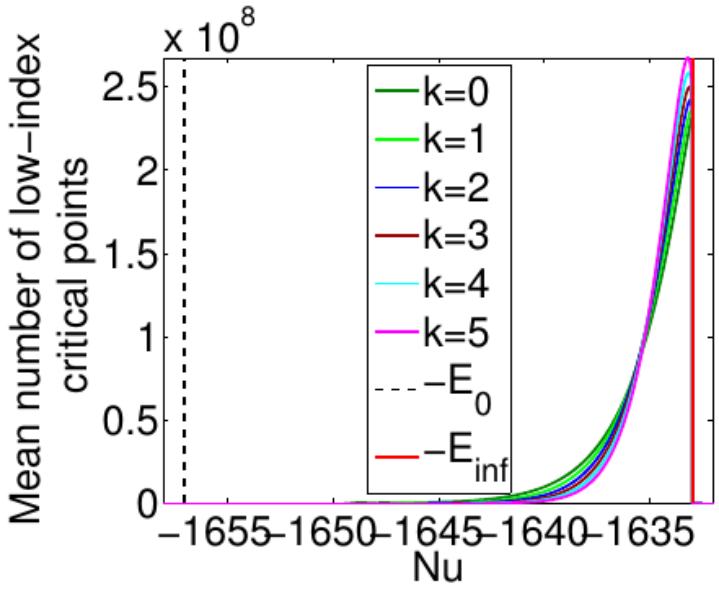

Zoomed:

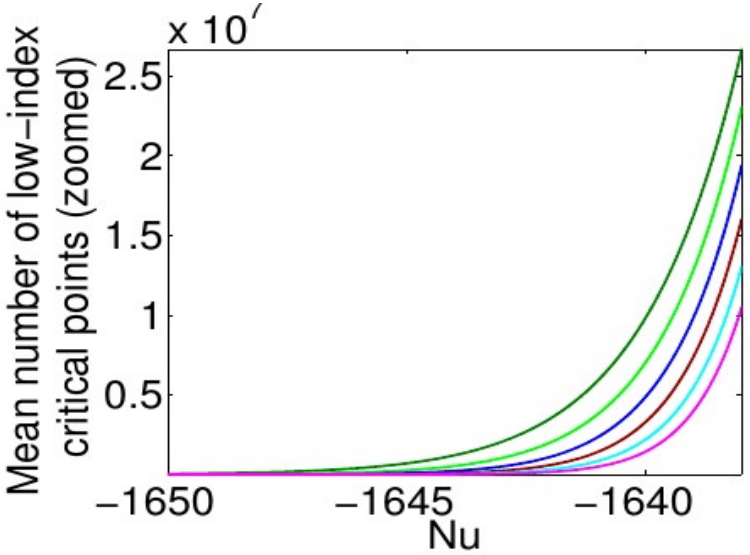

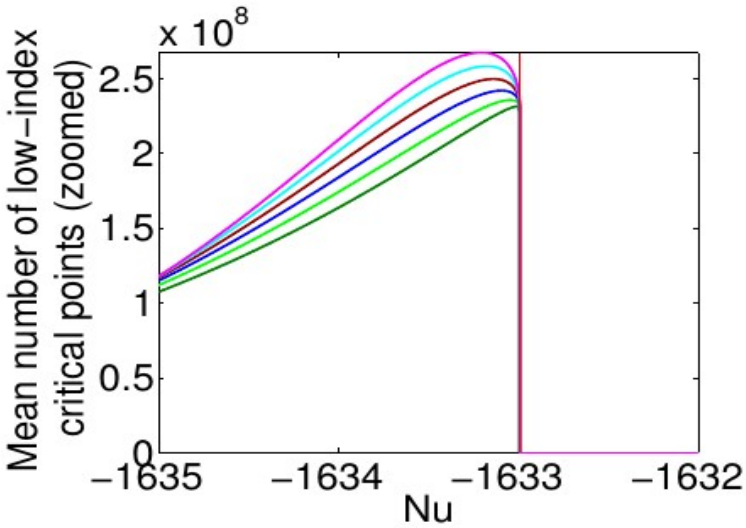

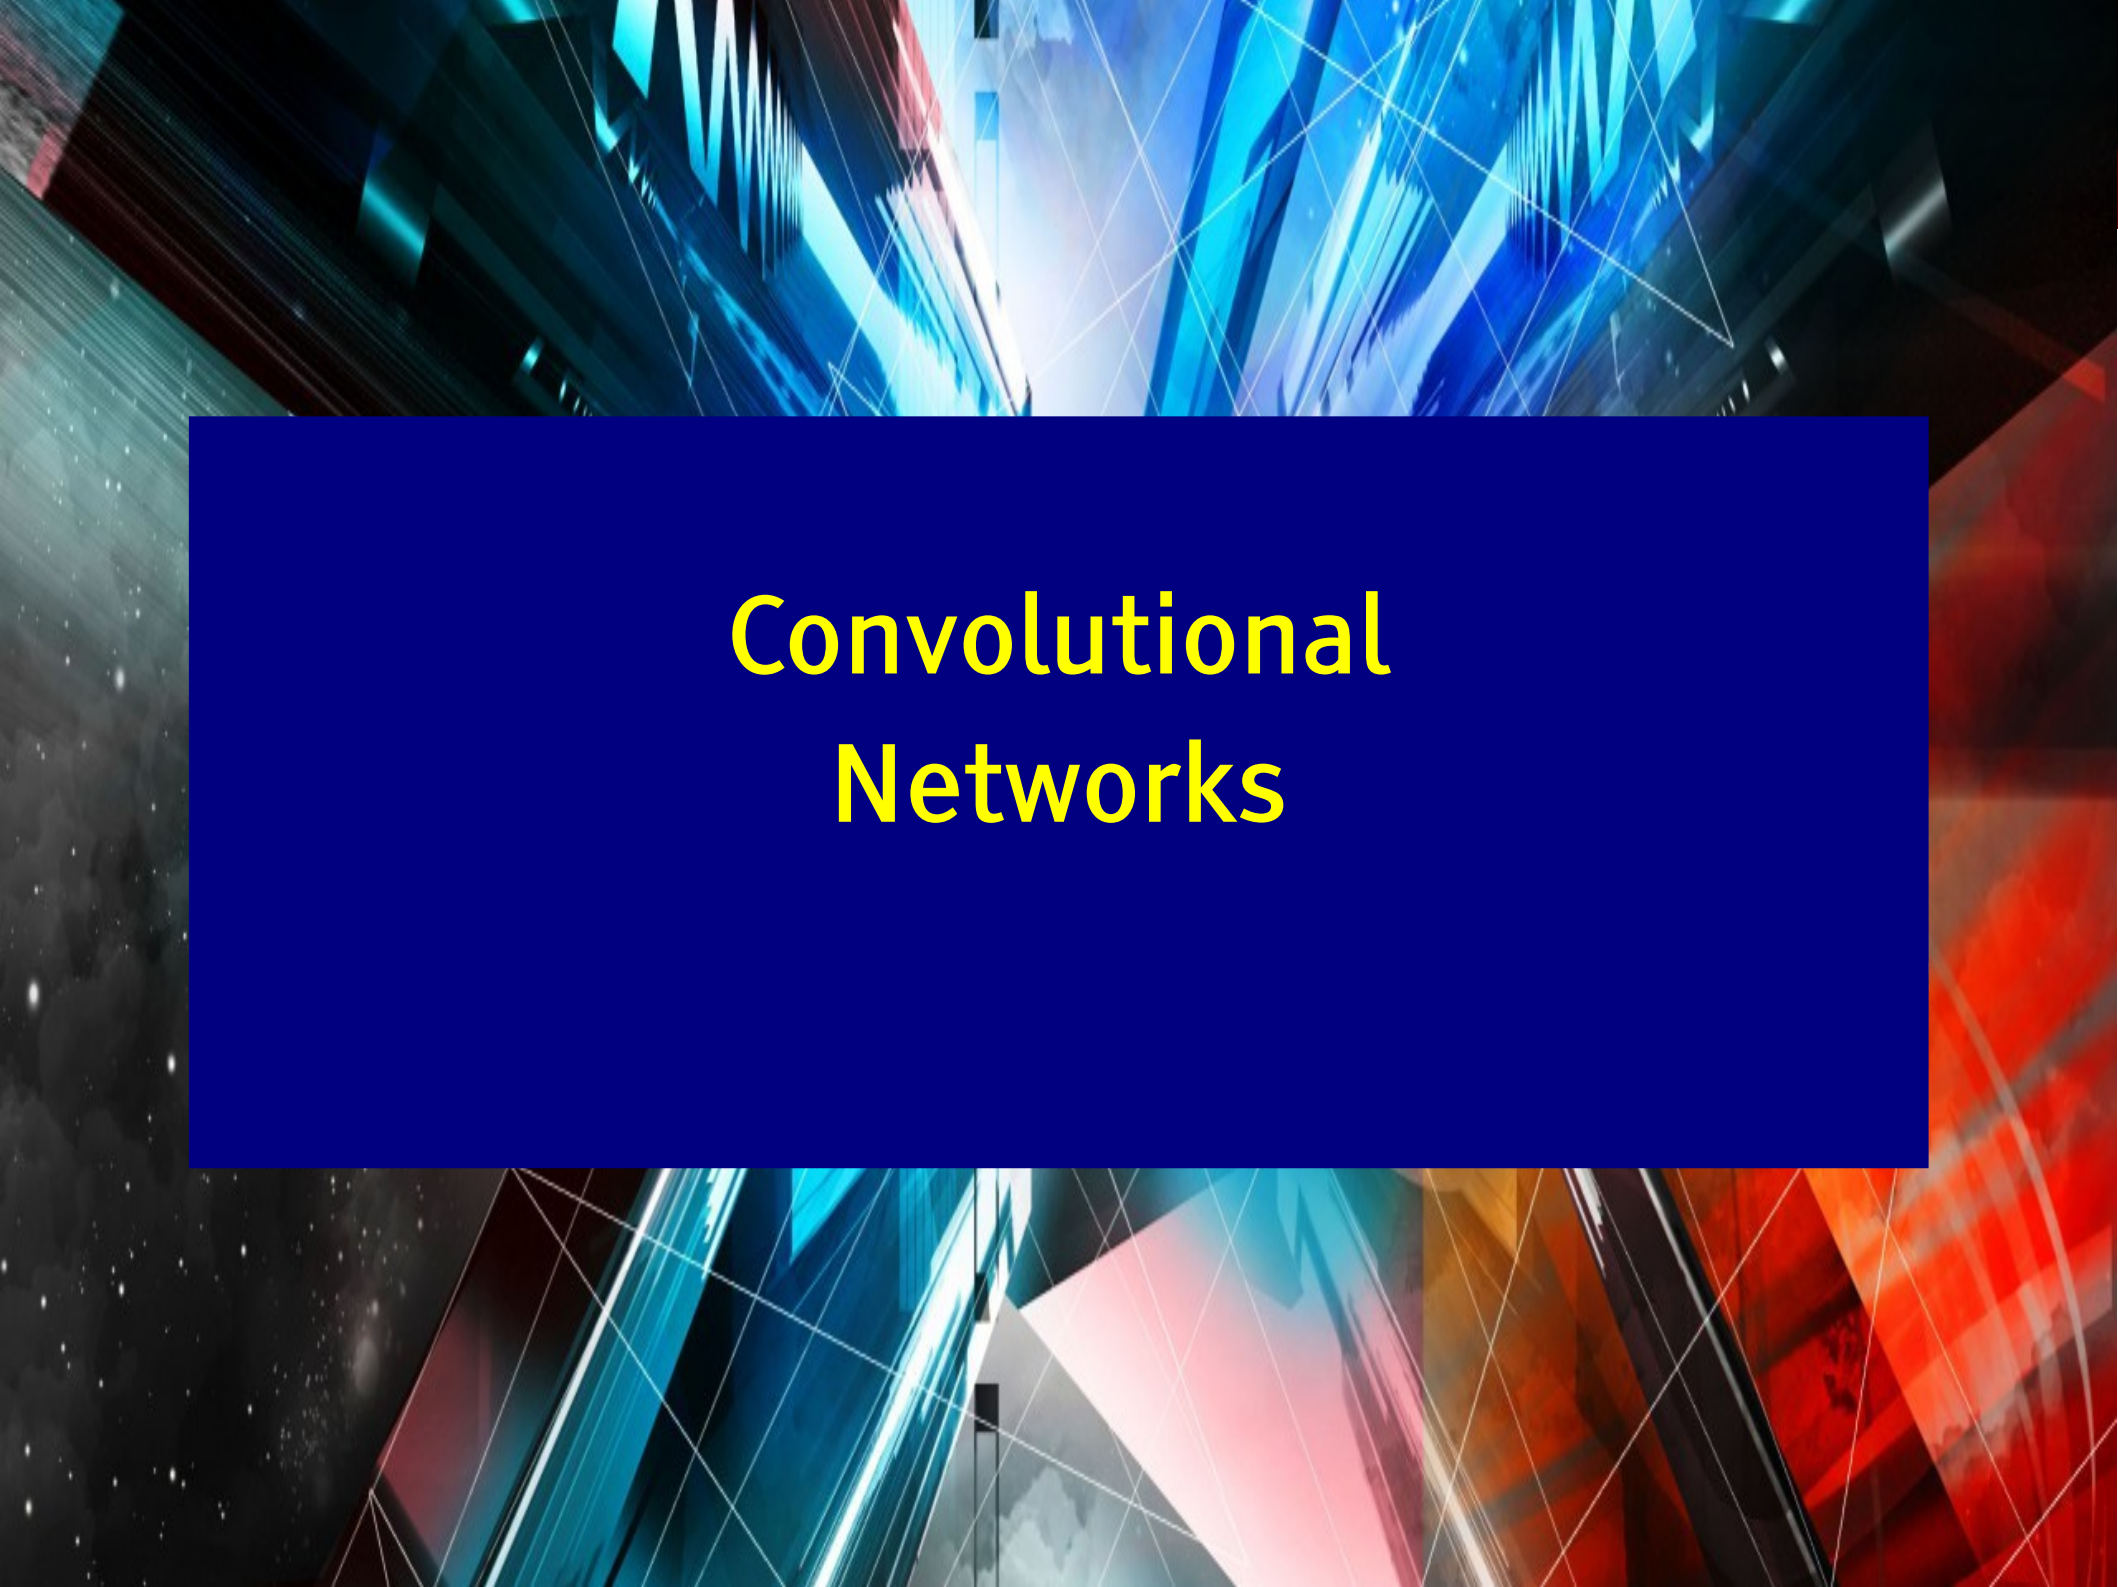
Convolutional
Networks

Convolutional Network

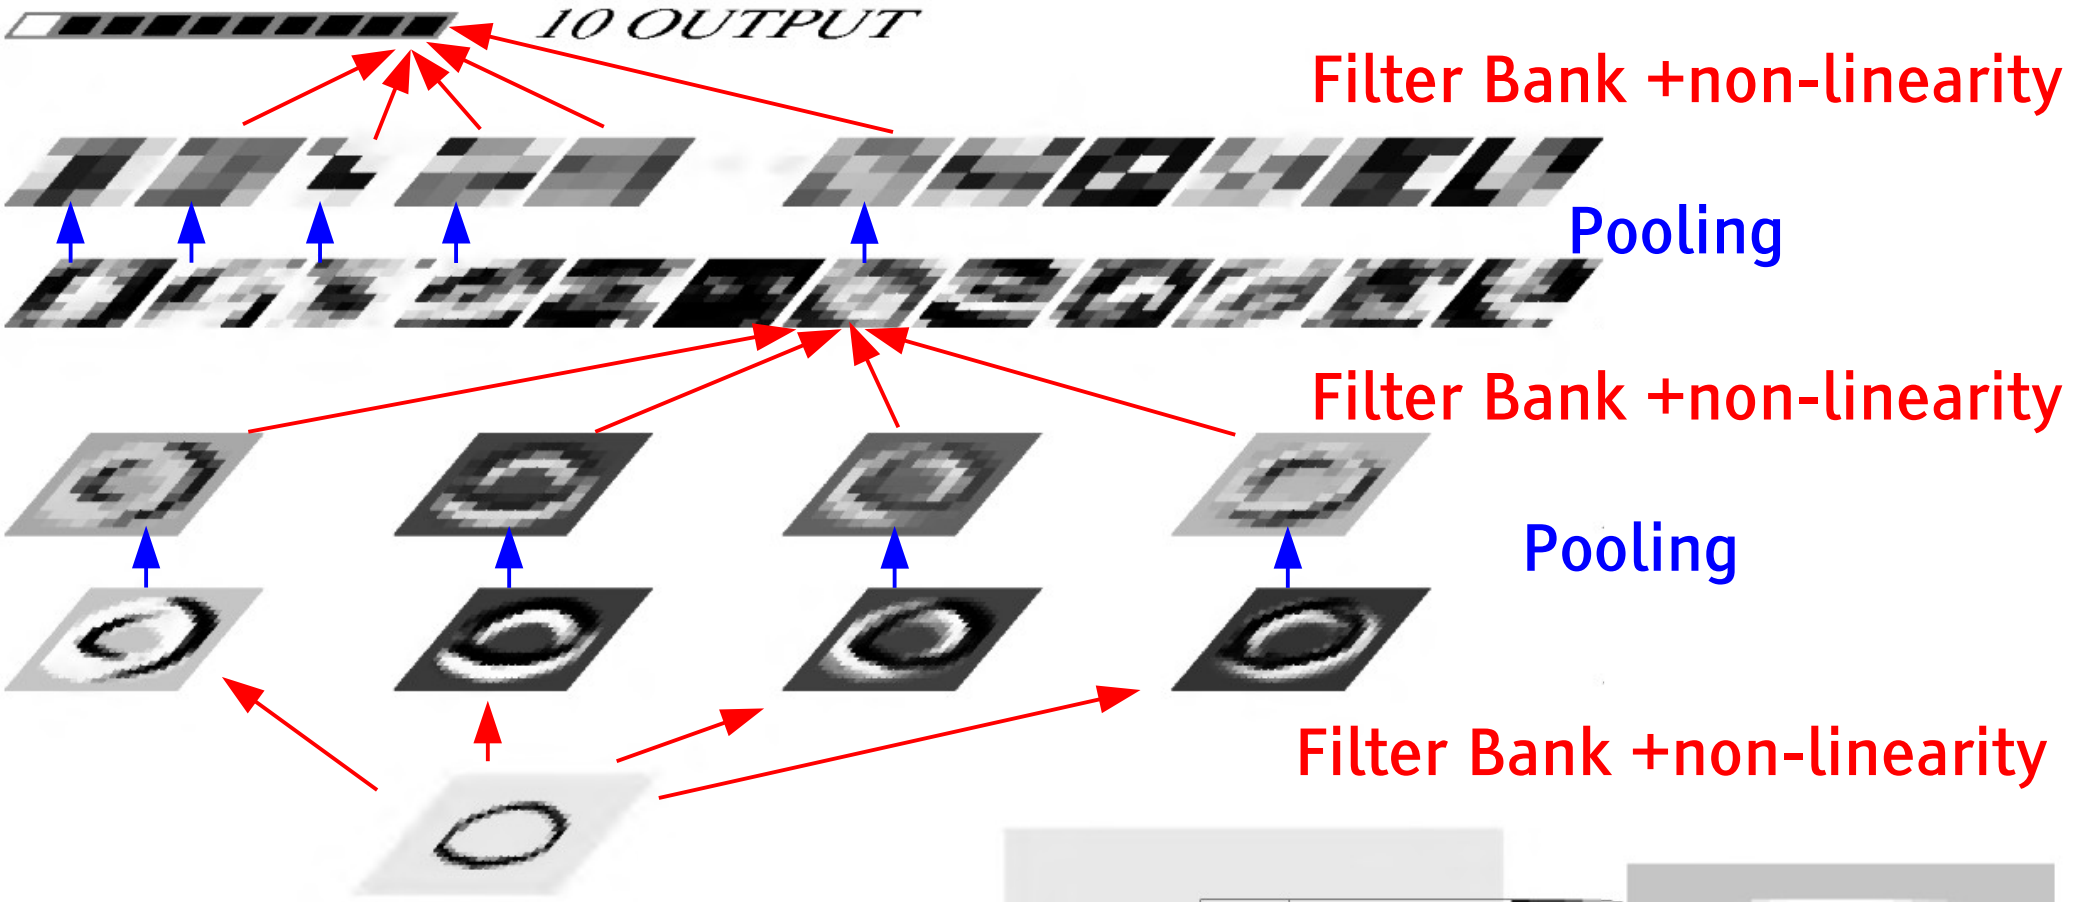

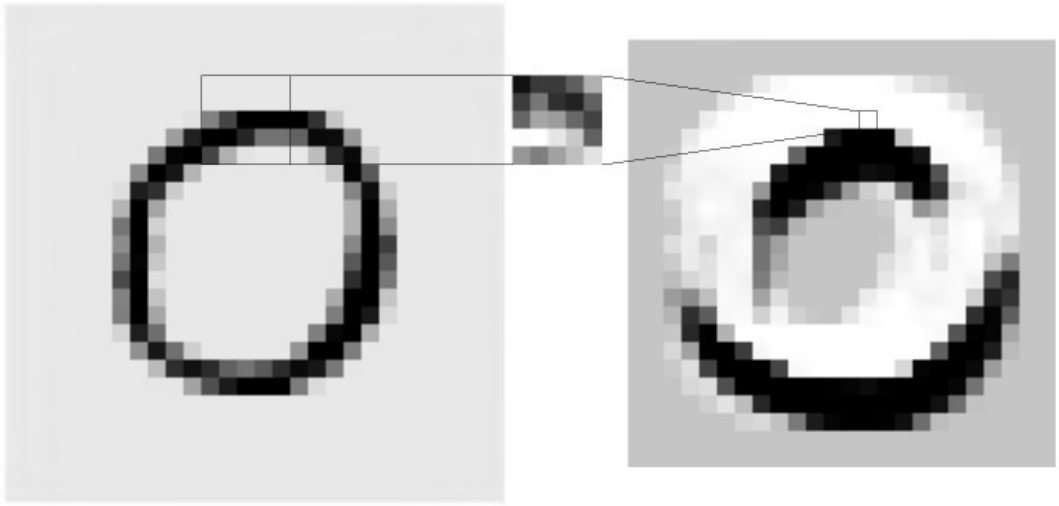

[LeCun et al. NIPS 1989]

Early Hierarchical Feature Models for Vision

- [Hubel & Wiesel 1962]:
  - simple cells detect local features
  - complex cells "pool" the outputs of simple cells within a retinotopic neighborhood.

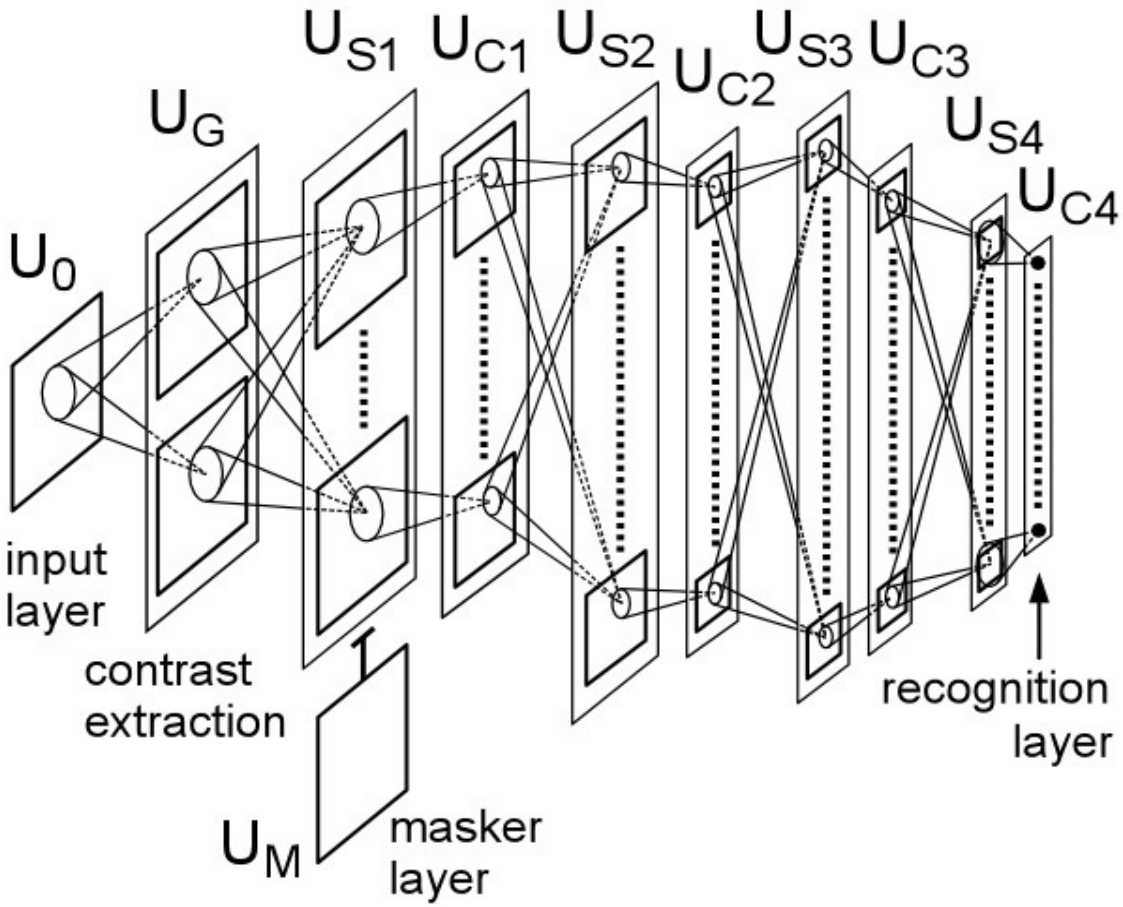

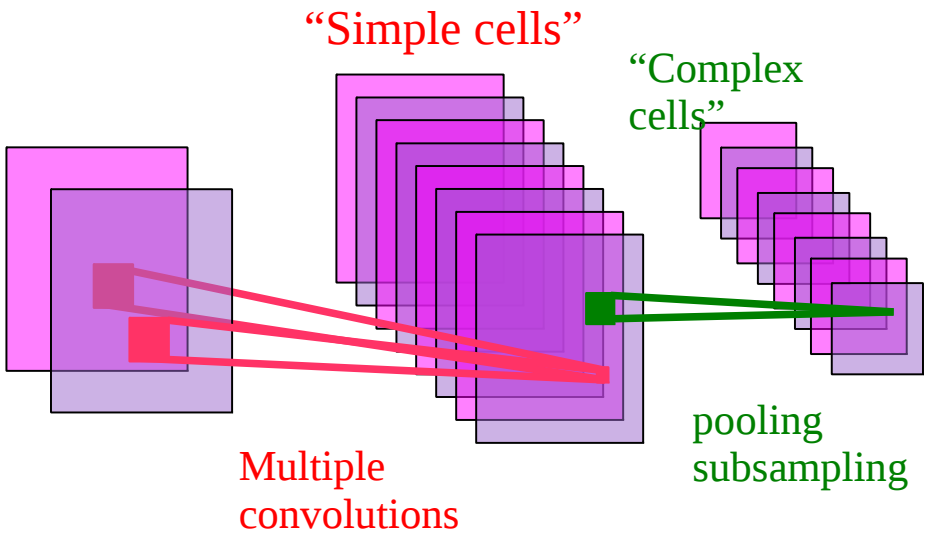

Cognitron & Neocognitron [Fukushima 1974-1982]

The Convolutional Net Model
(Multistage Hubel-Wiesel system)

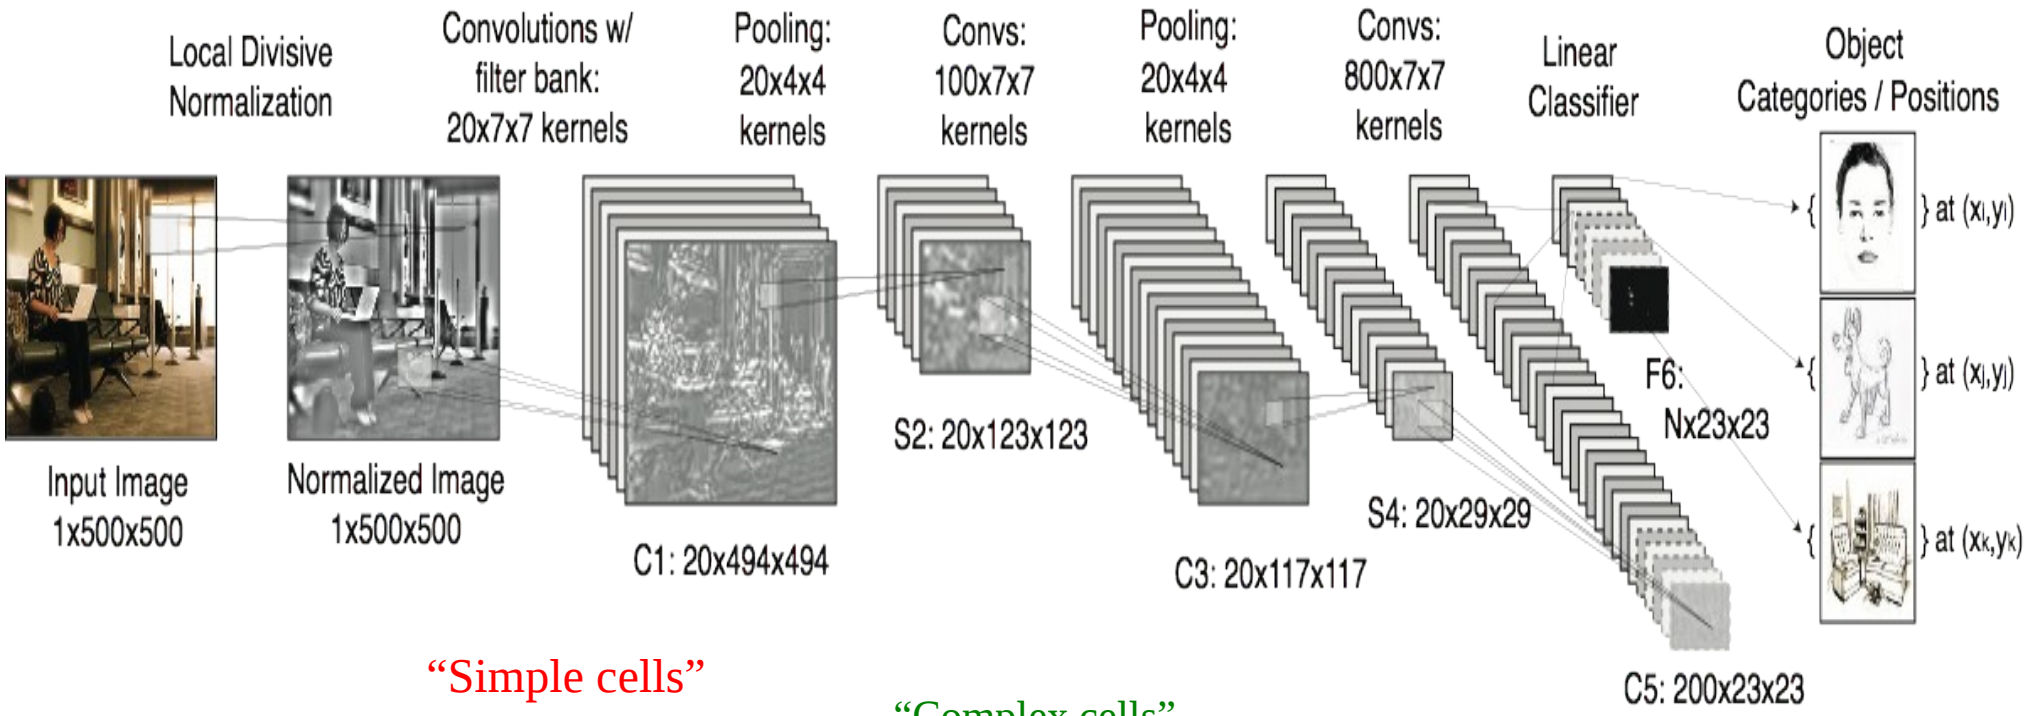

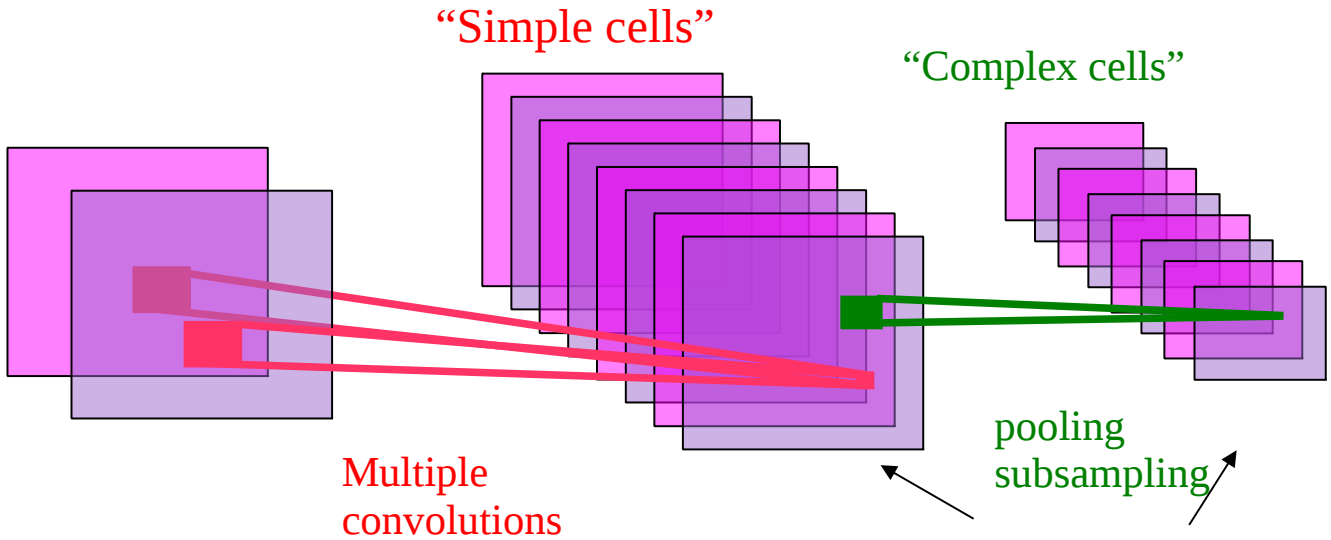

- Training is supervised
- With stochastic gradient descent

[LeCun et al. 89]
[LeCun et al. 98]

Convolutional Network (ConvNet)

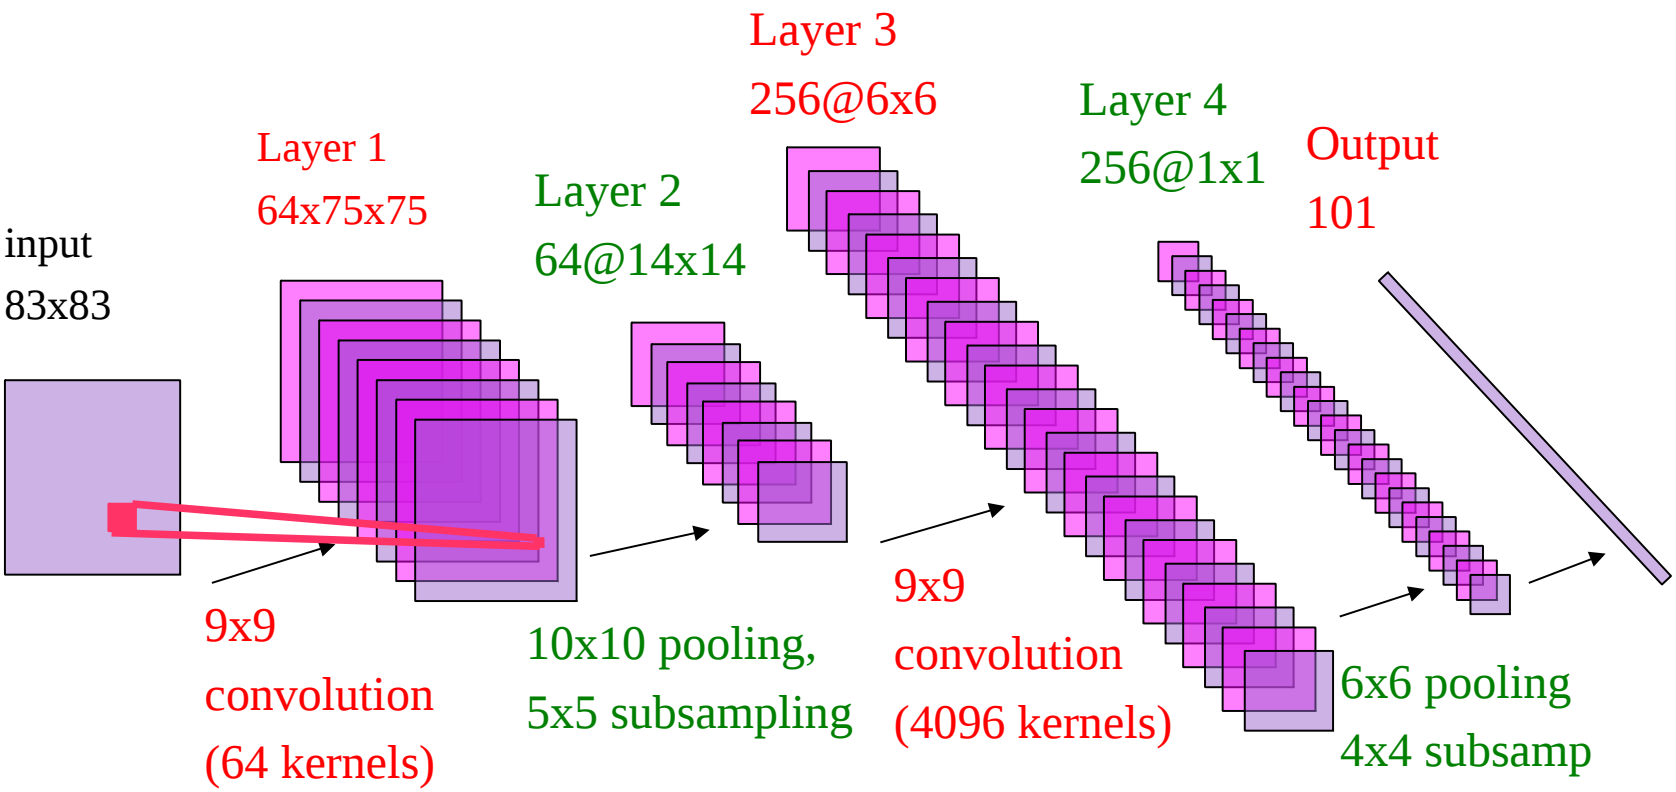

- Non-Linearity: half-wave rectification (ReLU), shrinkage function, sigmoid
- Pooling: max, average, L1, L2, log-sum-exp
- Training: Supervised (1988-2006), Unsupervised+Supervised (2006-now)

Convolutional Network (vintage 1990)

- filters → tanh → average-tanh → filters → tanh → average-tanh → filters → tanh

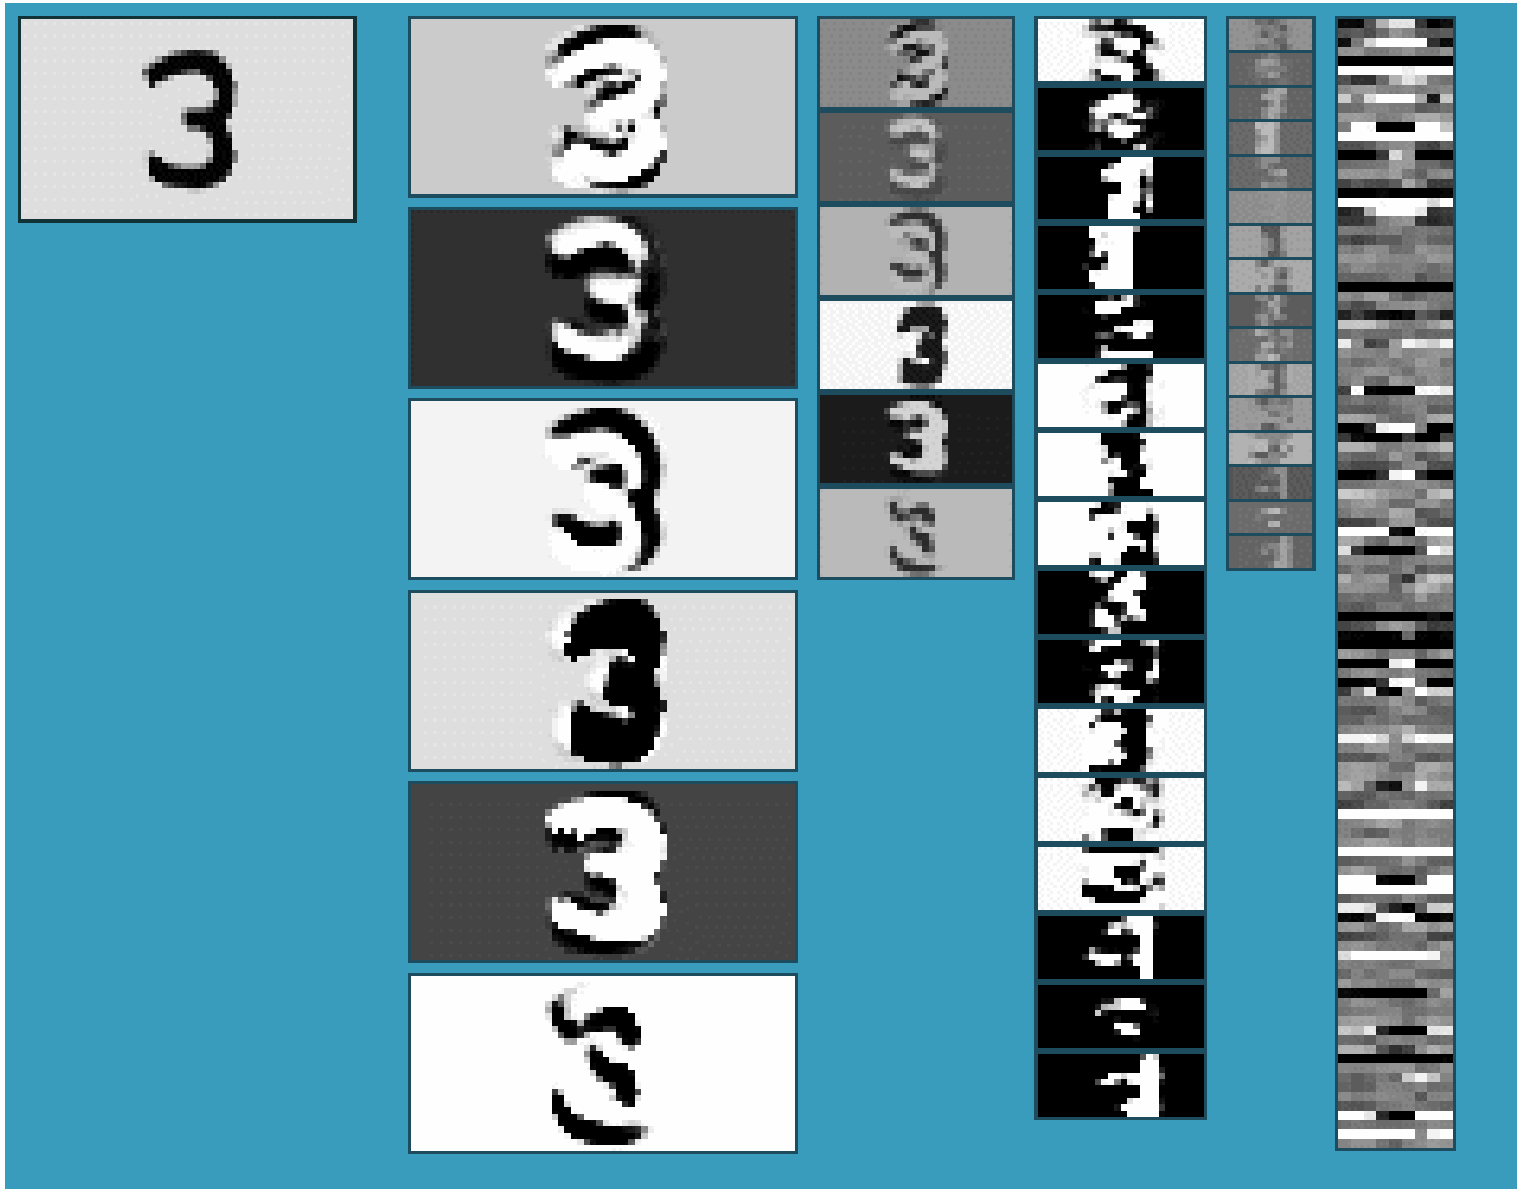

Curved manifold

Flatter manifold

LeNet1 Demo from 1993

- Running on a 486 PC with an AT&T DSP32C add-on board (20 Mflops!)

"Mainstream" object recognition pipeline 2006-2012: somewhat similar to ConvNets

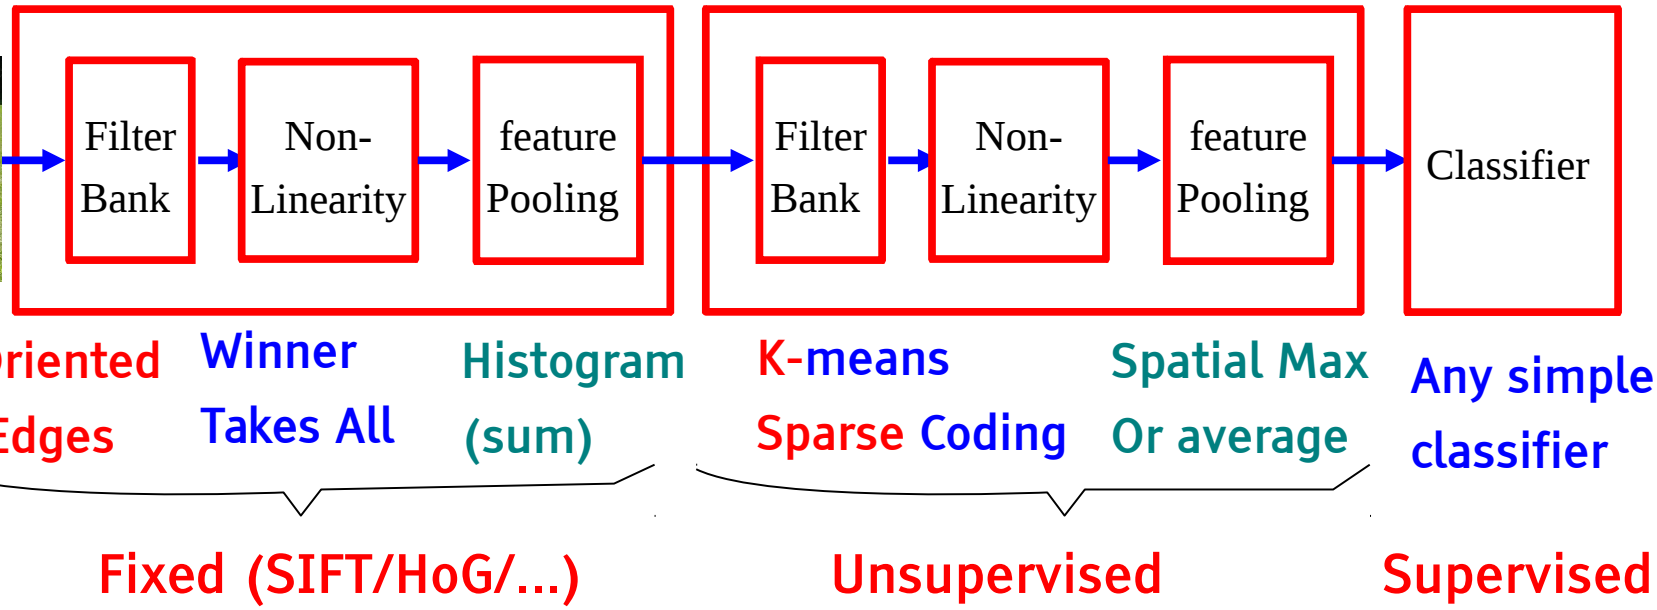

- Fixed Features + unsupervised mid-level features + simple classifier
  - SIFT + Vector Quantization + Pyramid pooling + SVM
    - [Lazebnik et al. CVPR 2006]
  - SIFT + Local Sparse Coding Macrofeatures + Pyramid pooling + SVM
    - [Boureau et al. ICCV 2011]
  - SIFT + Fisher Vectors + Deformable Parts Pooling + SVM
    - [Perronin et al. 2012]

Brute Force Approach
To
Multiple Object Recognition

Idea #1: Sliding Window ConvNet + Weighted FSM

- "Space Displacement Neural Net".
- Convolutions are applied to a large image
- Output and feature maps are extended/replicated accordingly

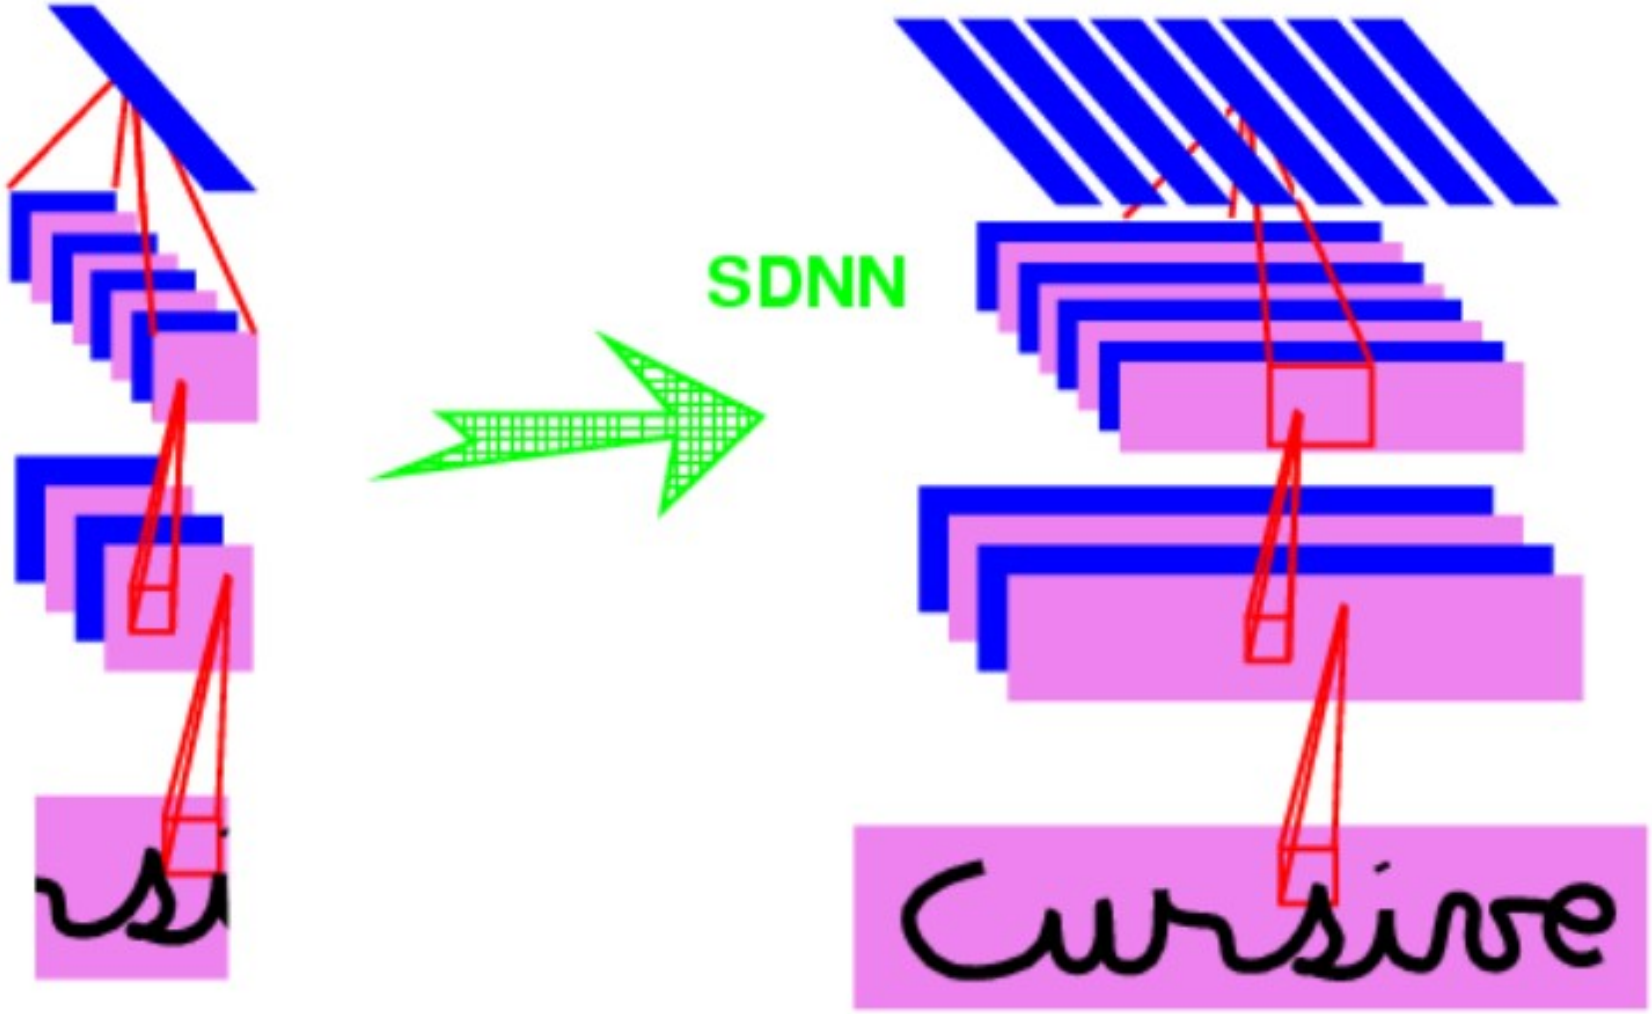

Idea #1: Sliding Window ConvNet + Weighted FSM

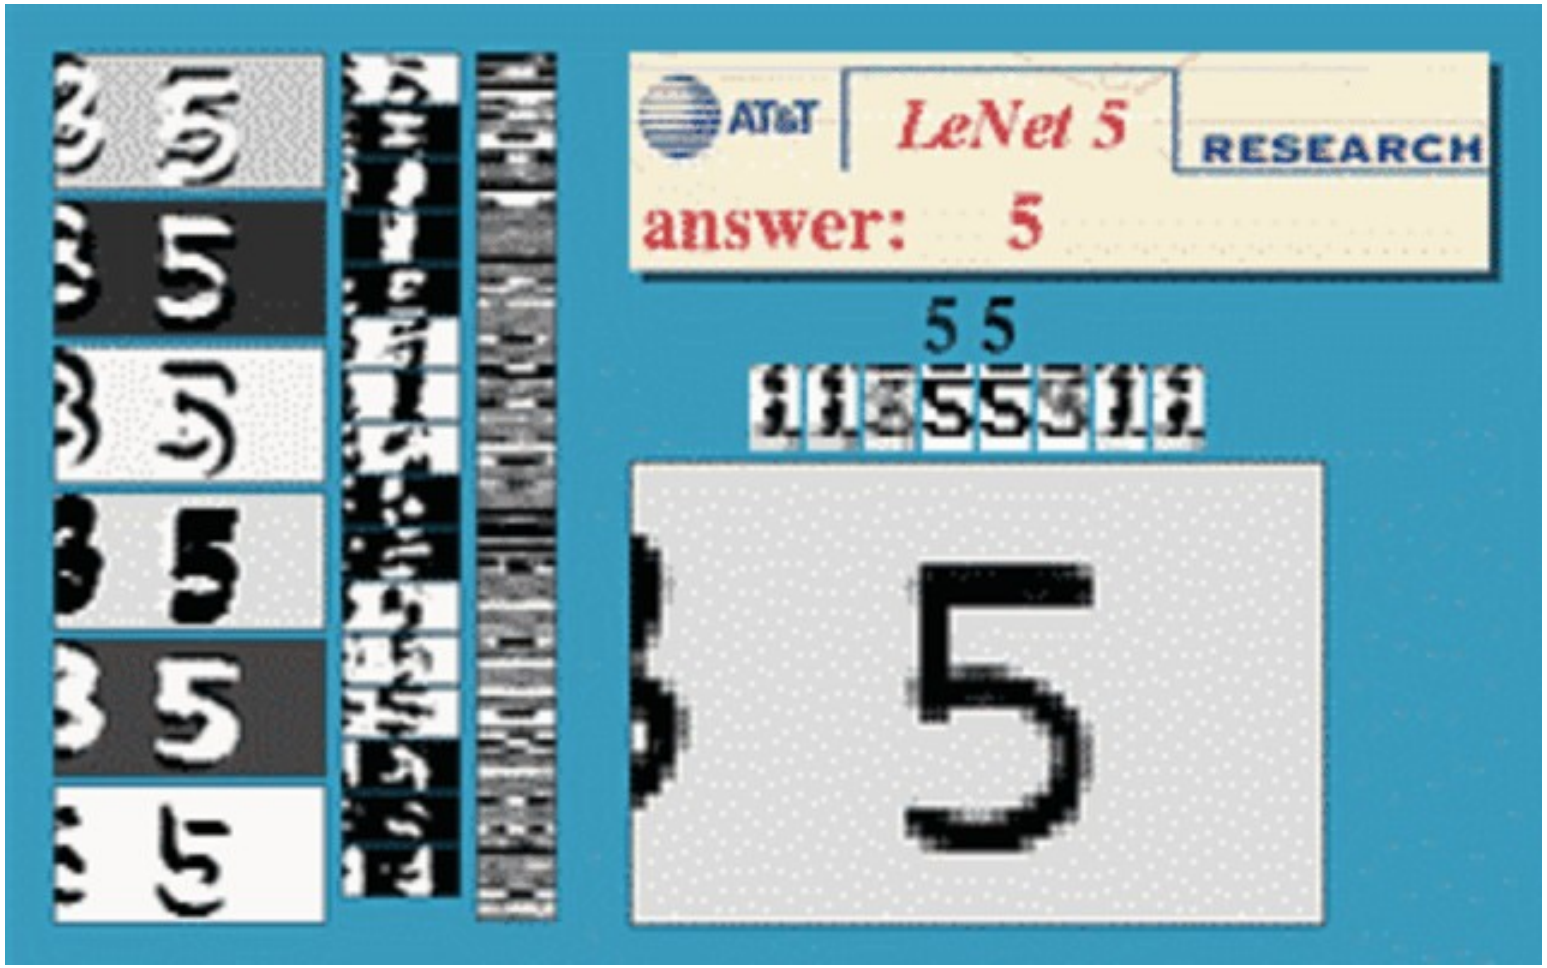

Idea #1: Sliding Window ConvNet + Weighted FSM

Idea #1: Sliding Window ConvNet + Weighted FSM

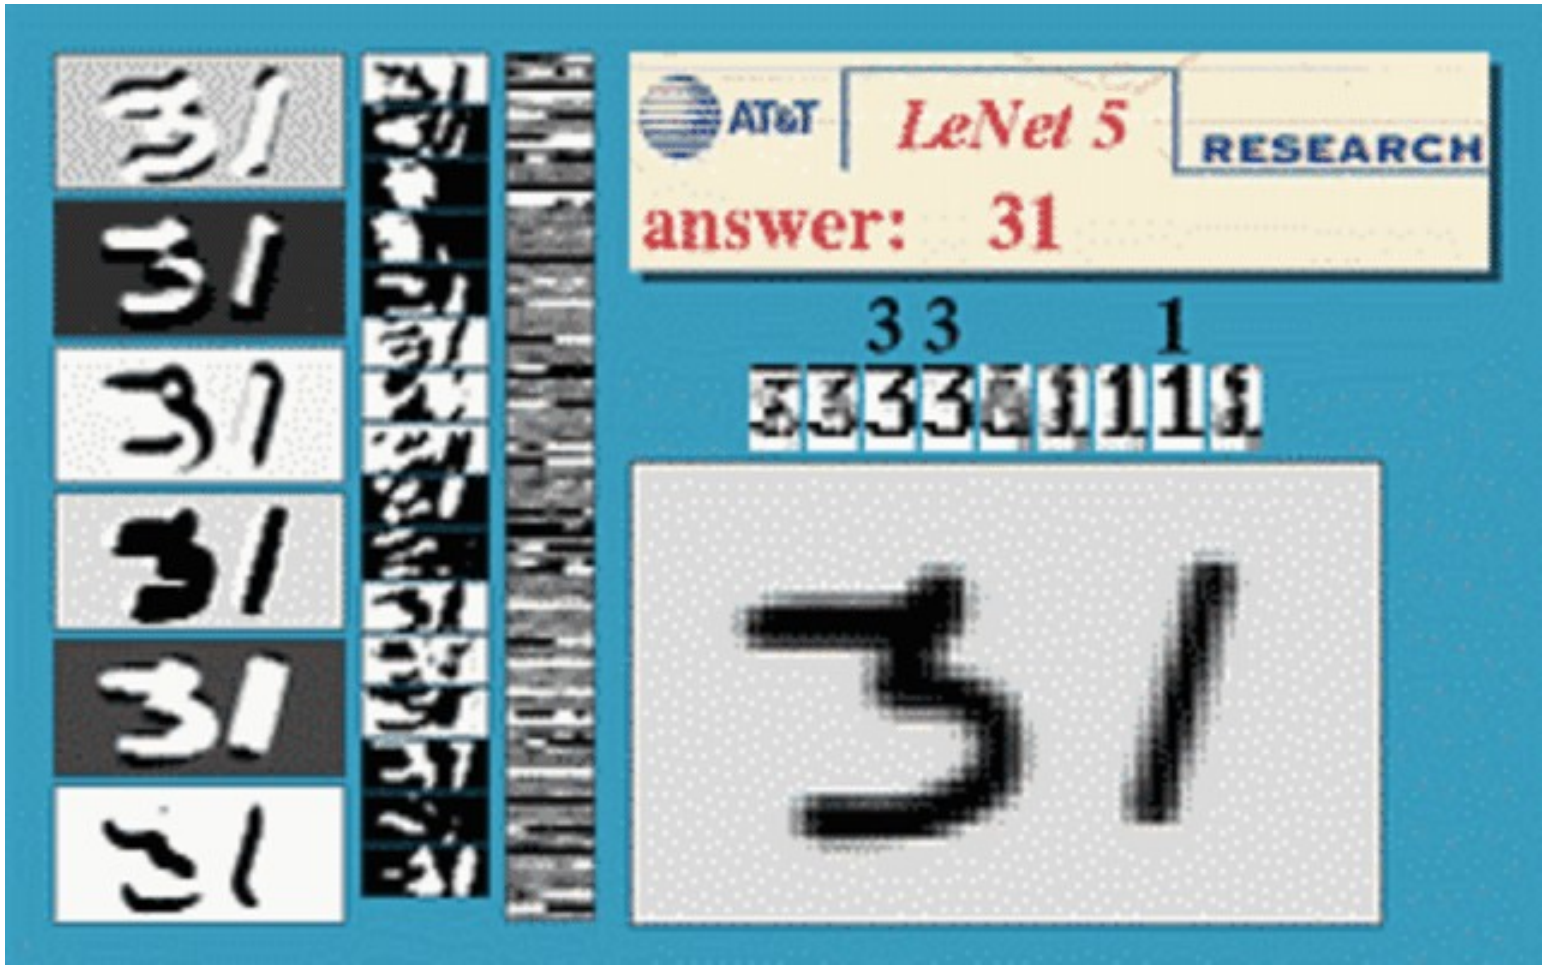

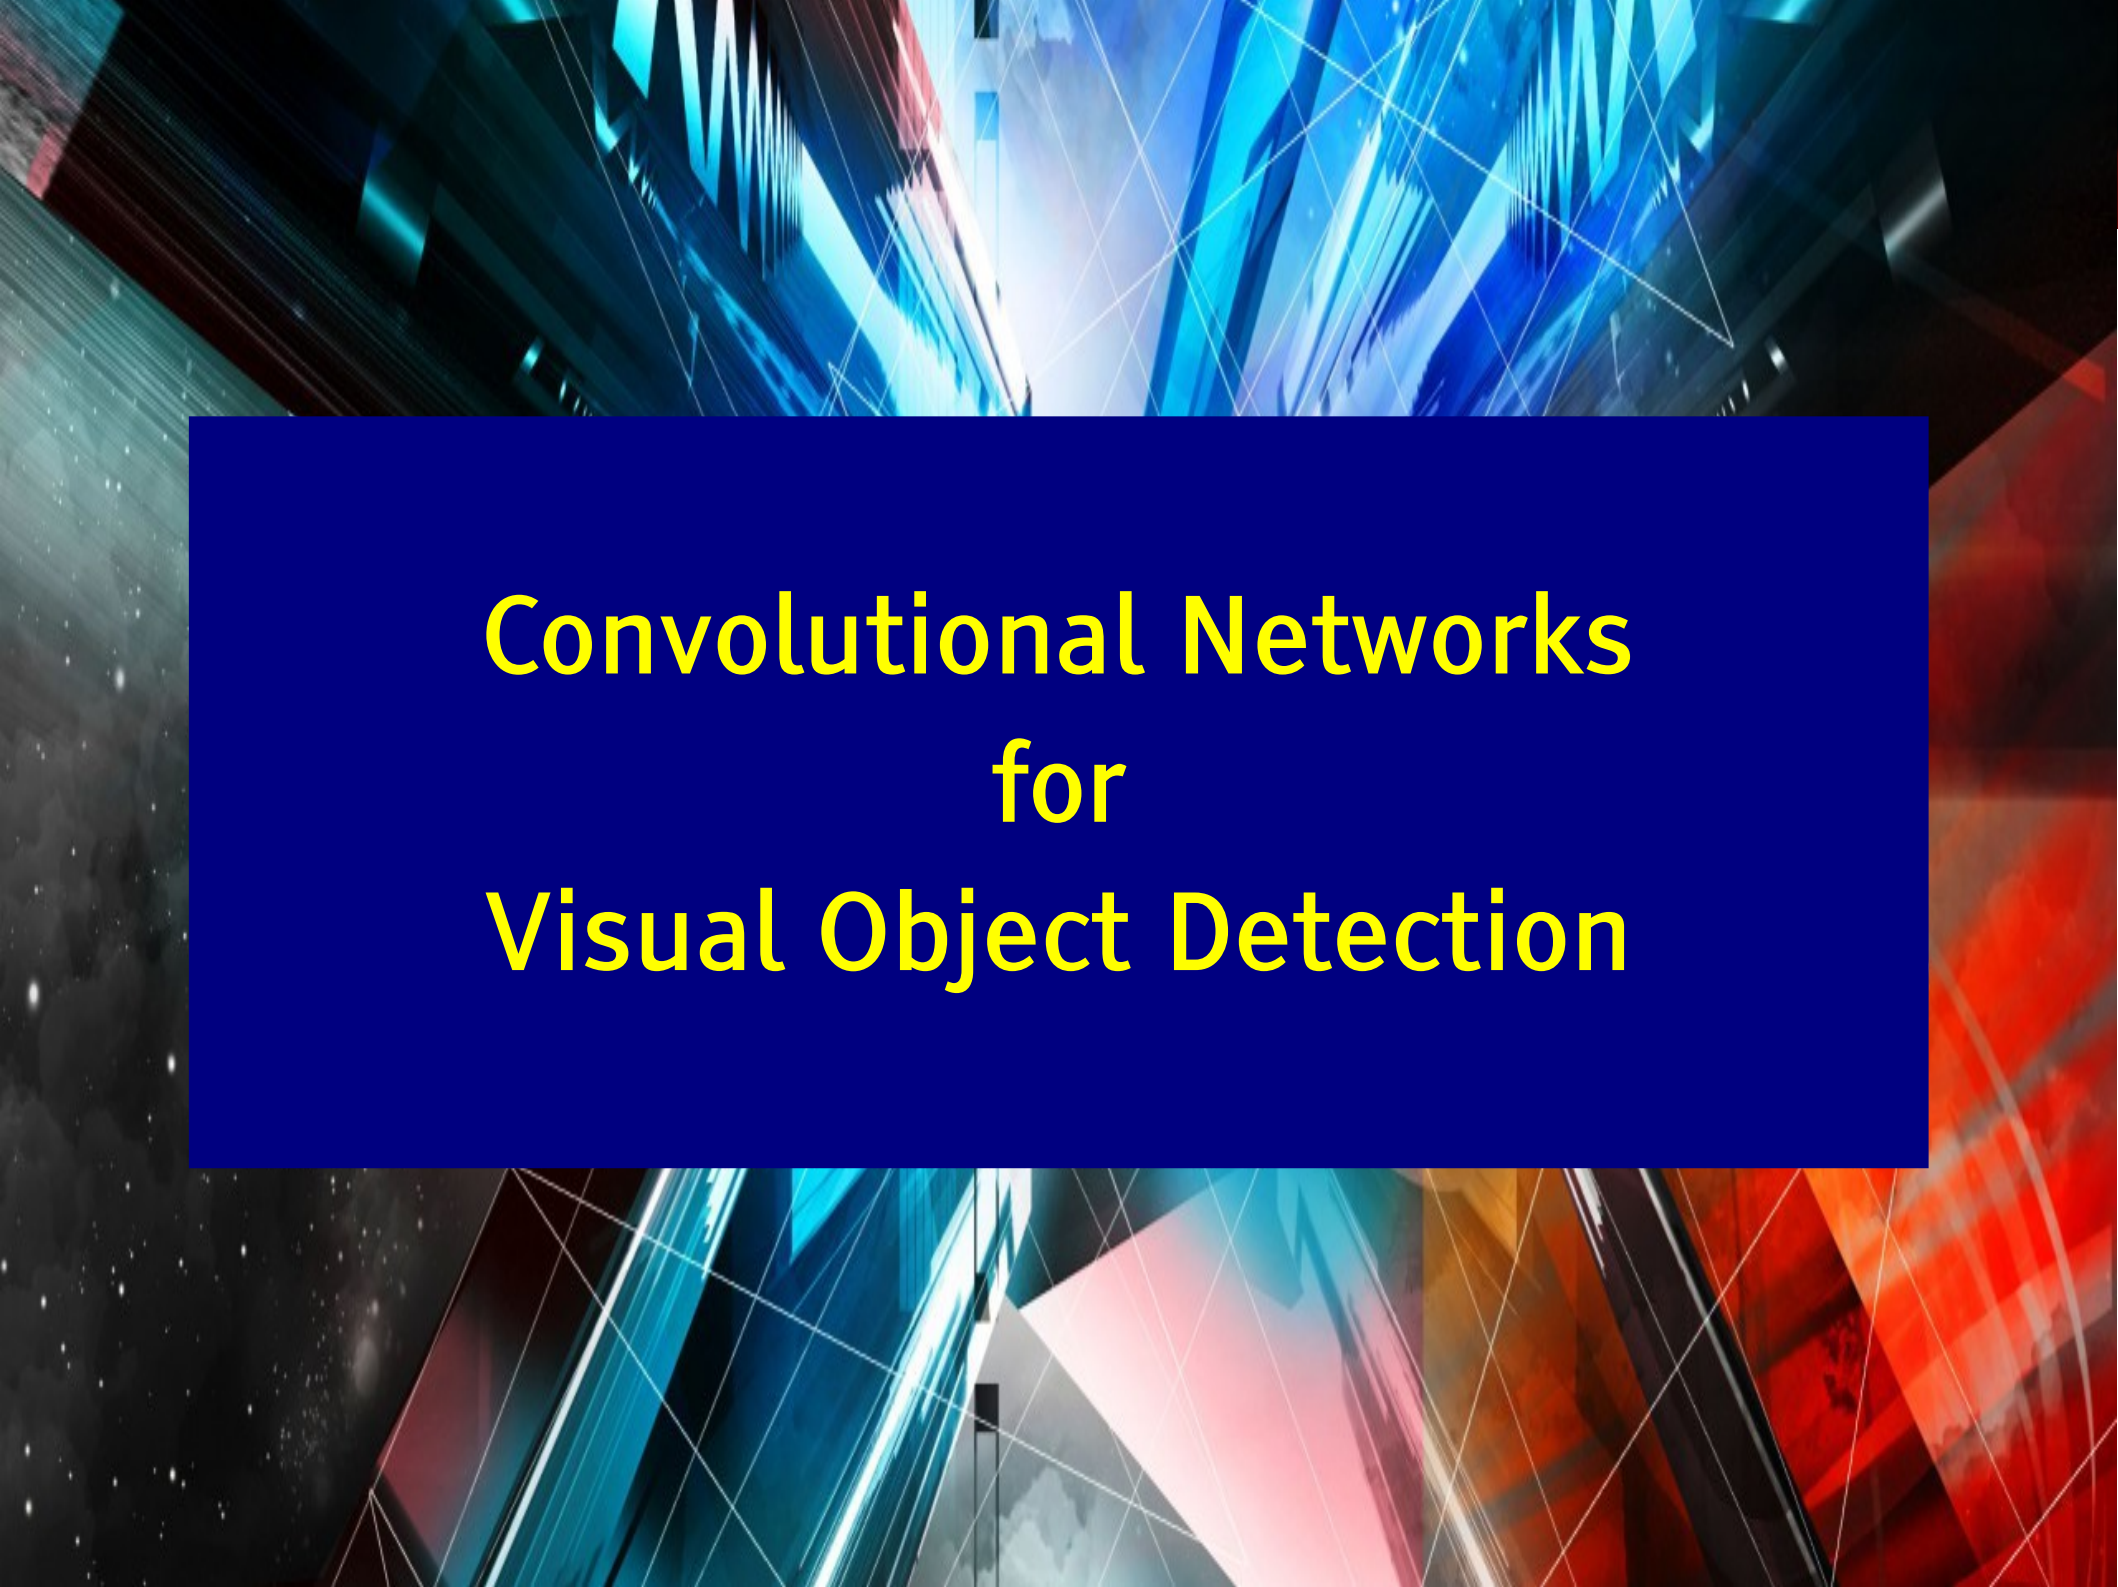
Convolutional Networks
for
Visual Object Detection

mid 2000s: state of the art results on face detection

Data Set->	TILTED		PROFILE		MIT+CMU	
False positives per image->	4.42	26.9	0.47	3.36	0.5	1.28
Our Detector	90%	97%	67%	83%	83%	88%
Jones & Viola (tilted)	90%	95%	x		x	
Jones & Viola (profile)	x		70%	83%	x	

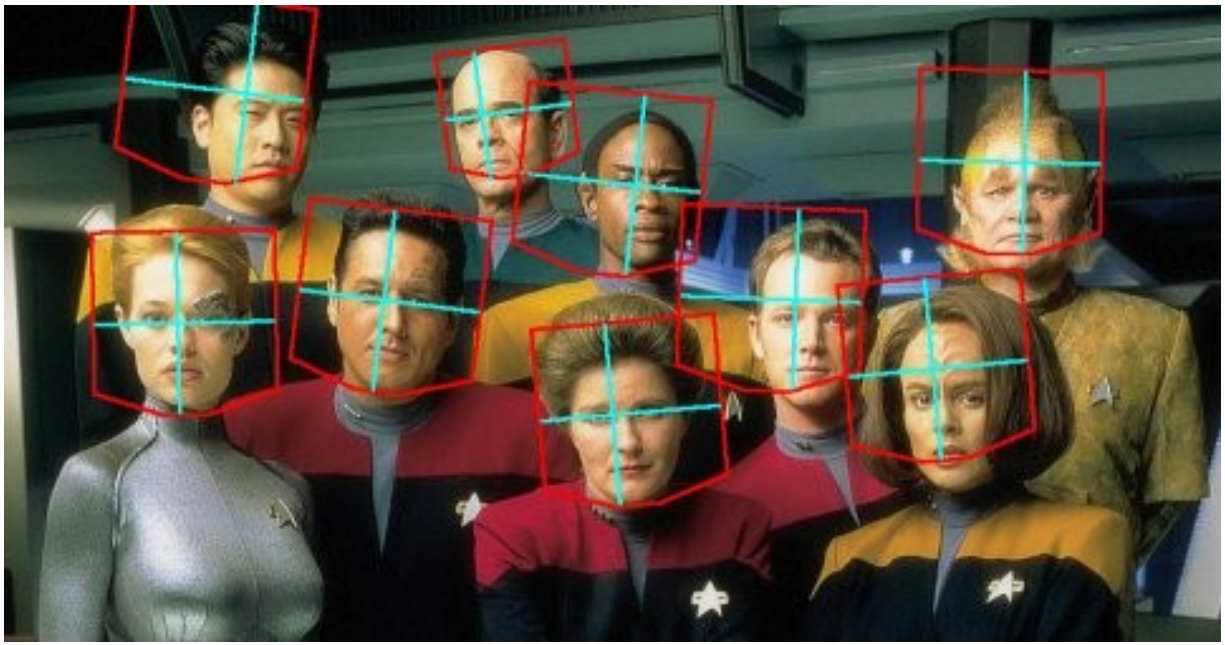

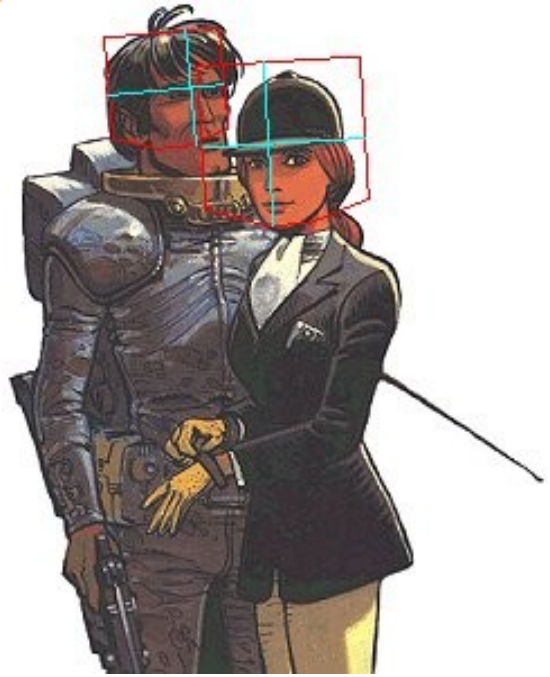

[Vaillant et al. IEE 1994] [Osadchy et al. 2004] [Osadchy et al, JMLR 2007]

Simultaneous face detection and pose estimation

Y LeCun

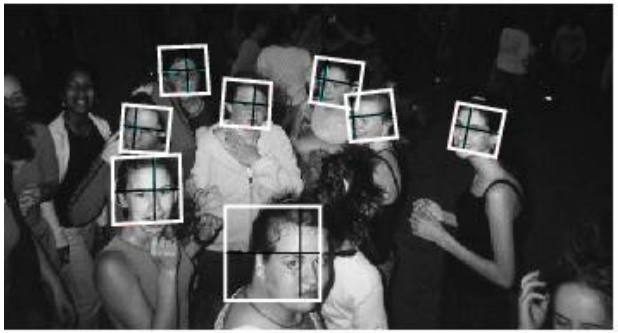

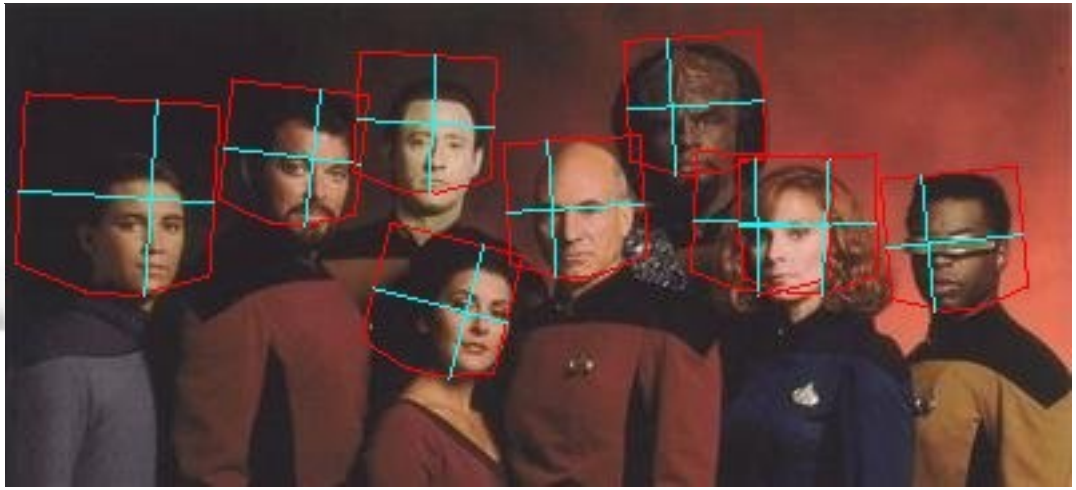
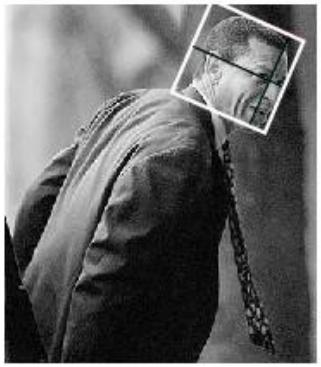

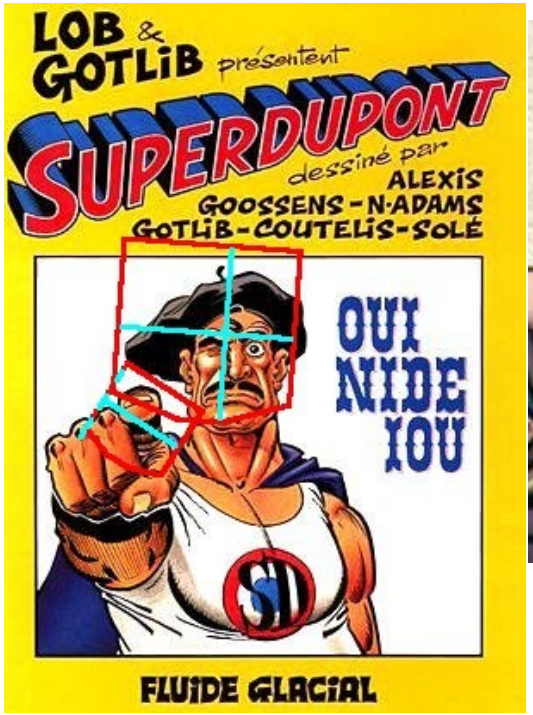

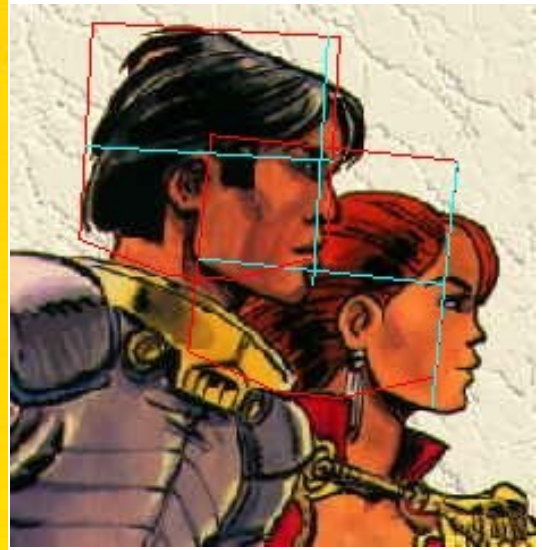

[Vaillant et al. IEE 1994] [Osadchy et al. 2004] [Osadchy et al, JMLR 2007]

Simultaneous face detection and pose estimation

Y LeCun

Pedestrian Detection: INRIA Dataset. Miss rate vs false positives

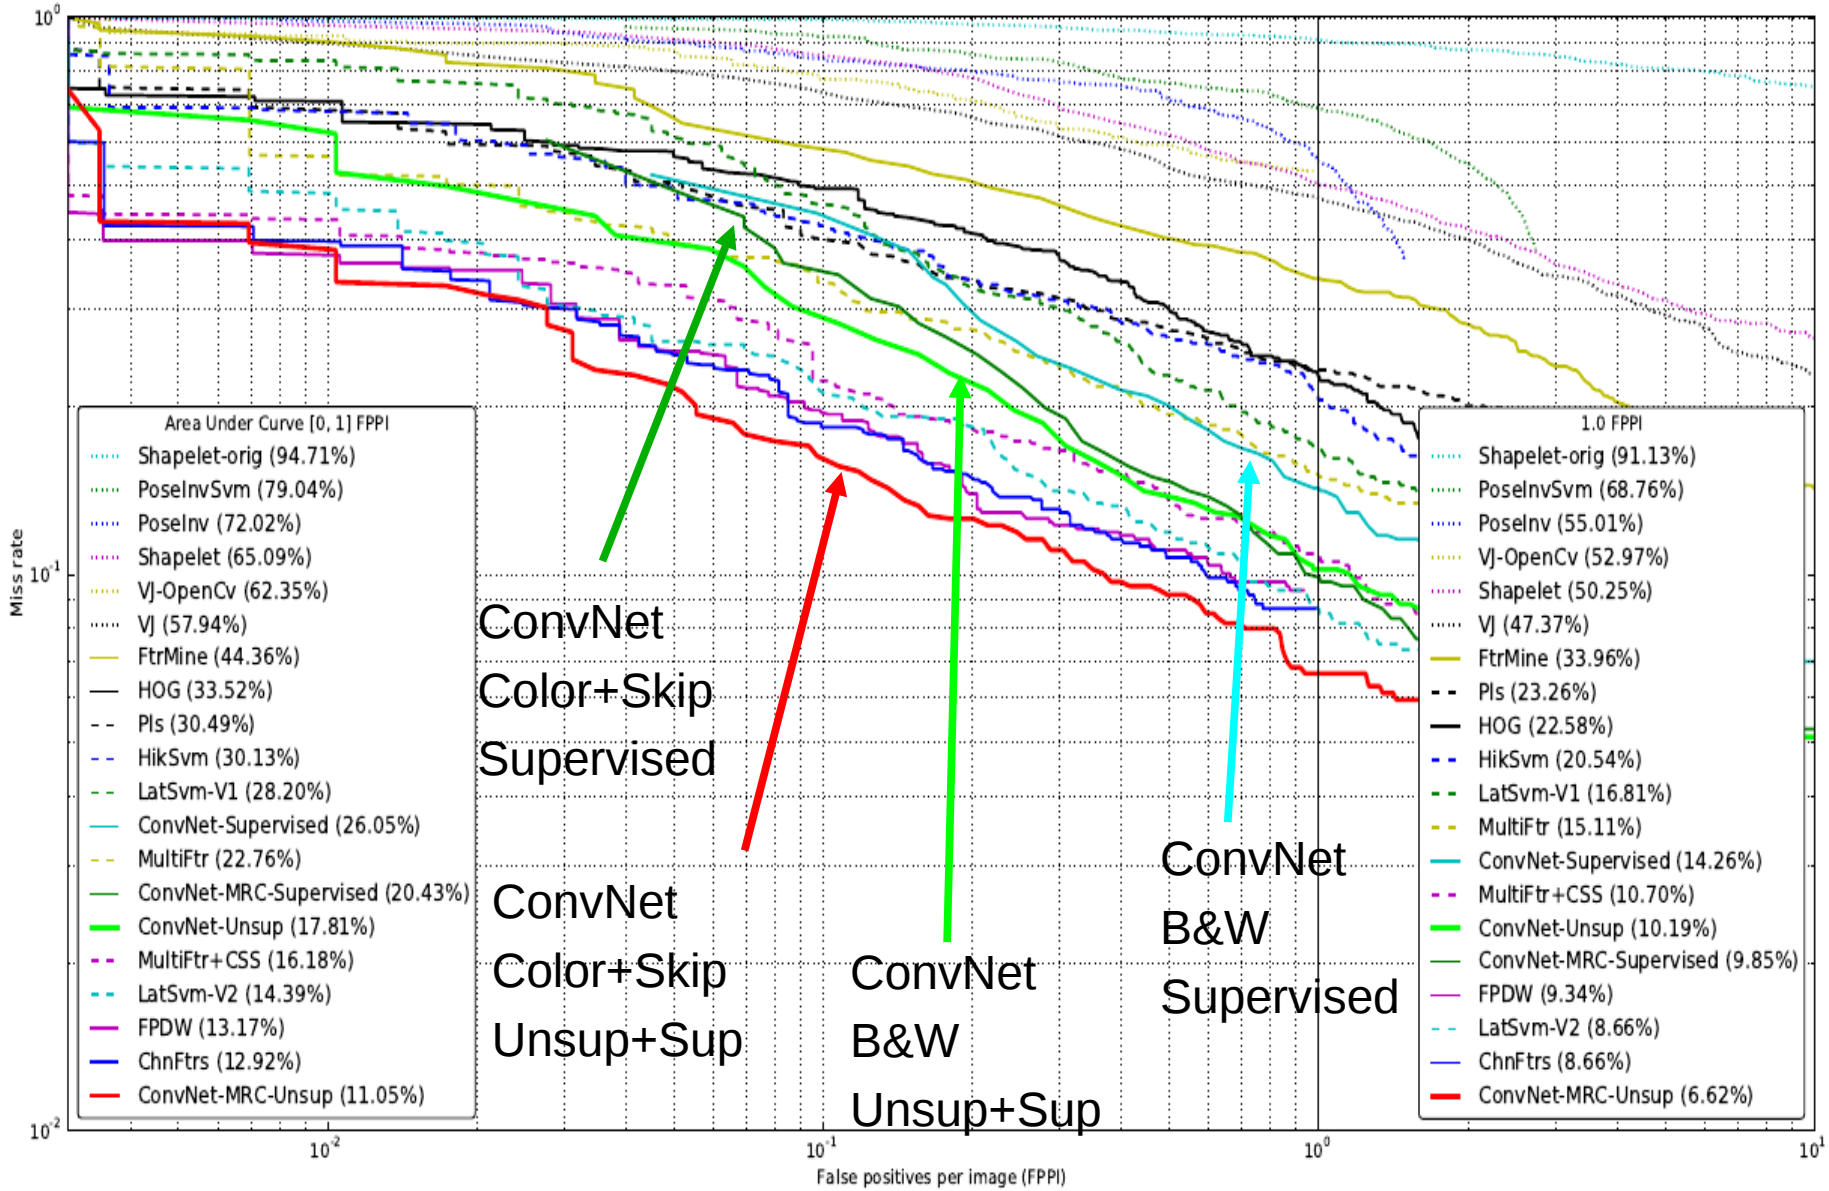

[Kavukcuoglu et al. NIPS 2010] [Sermanet et al. ArXiv 2012]

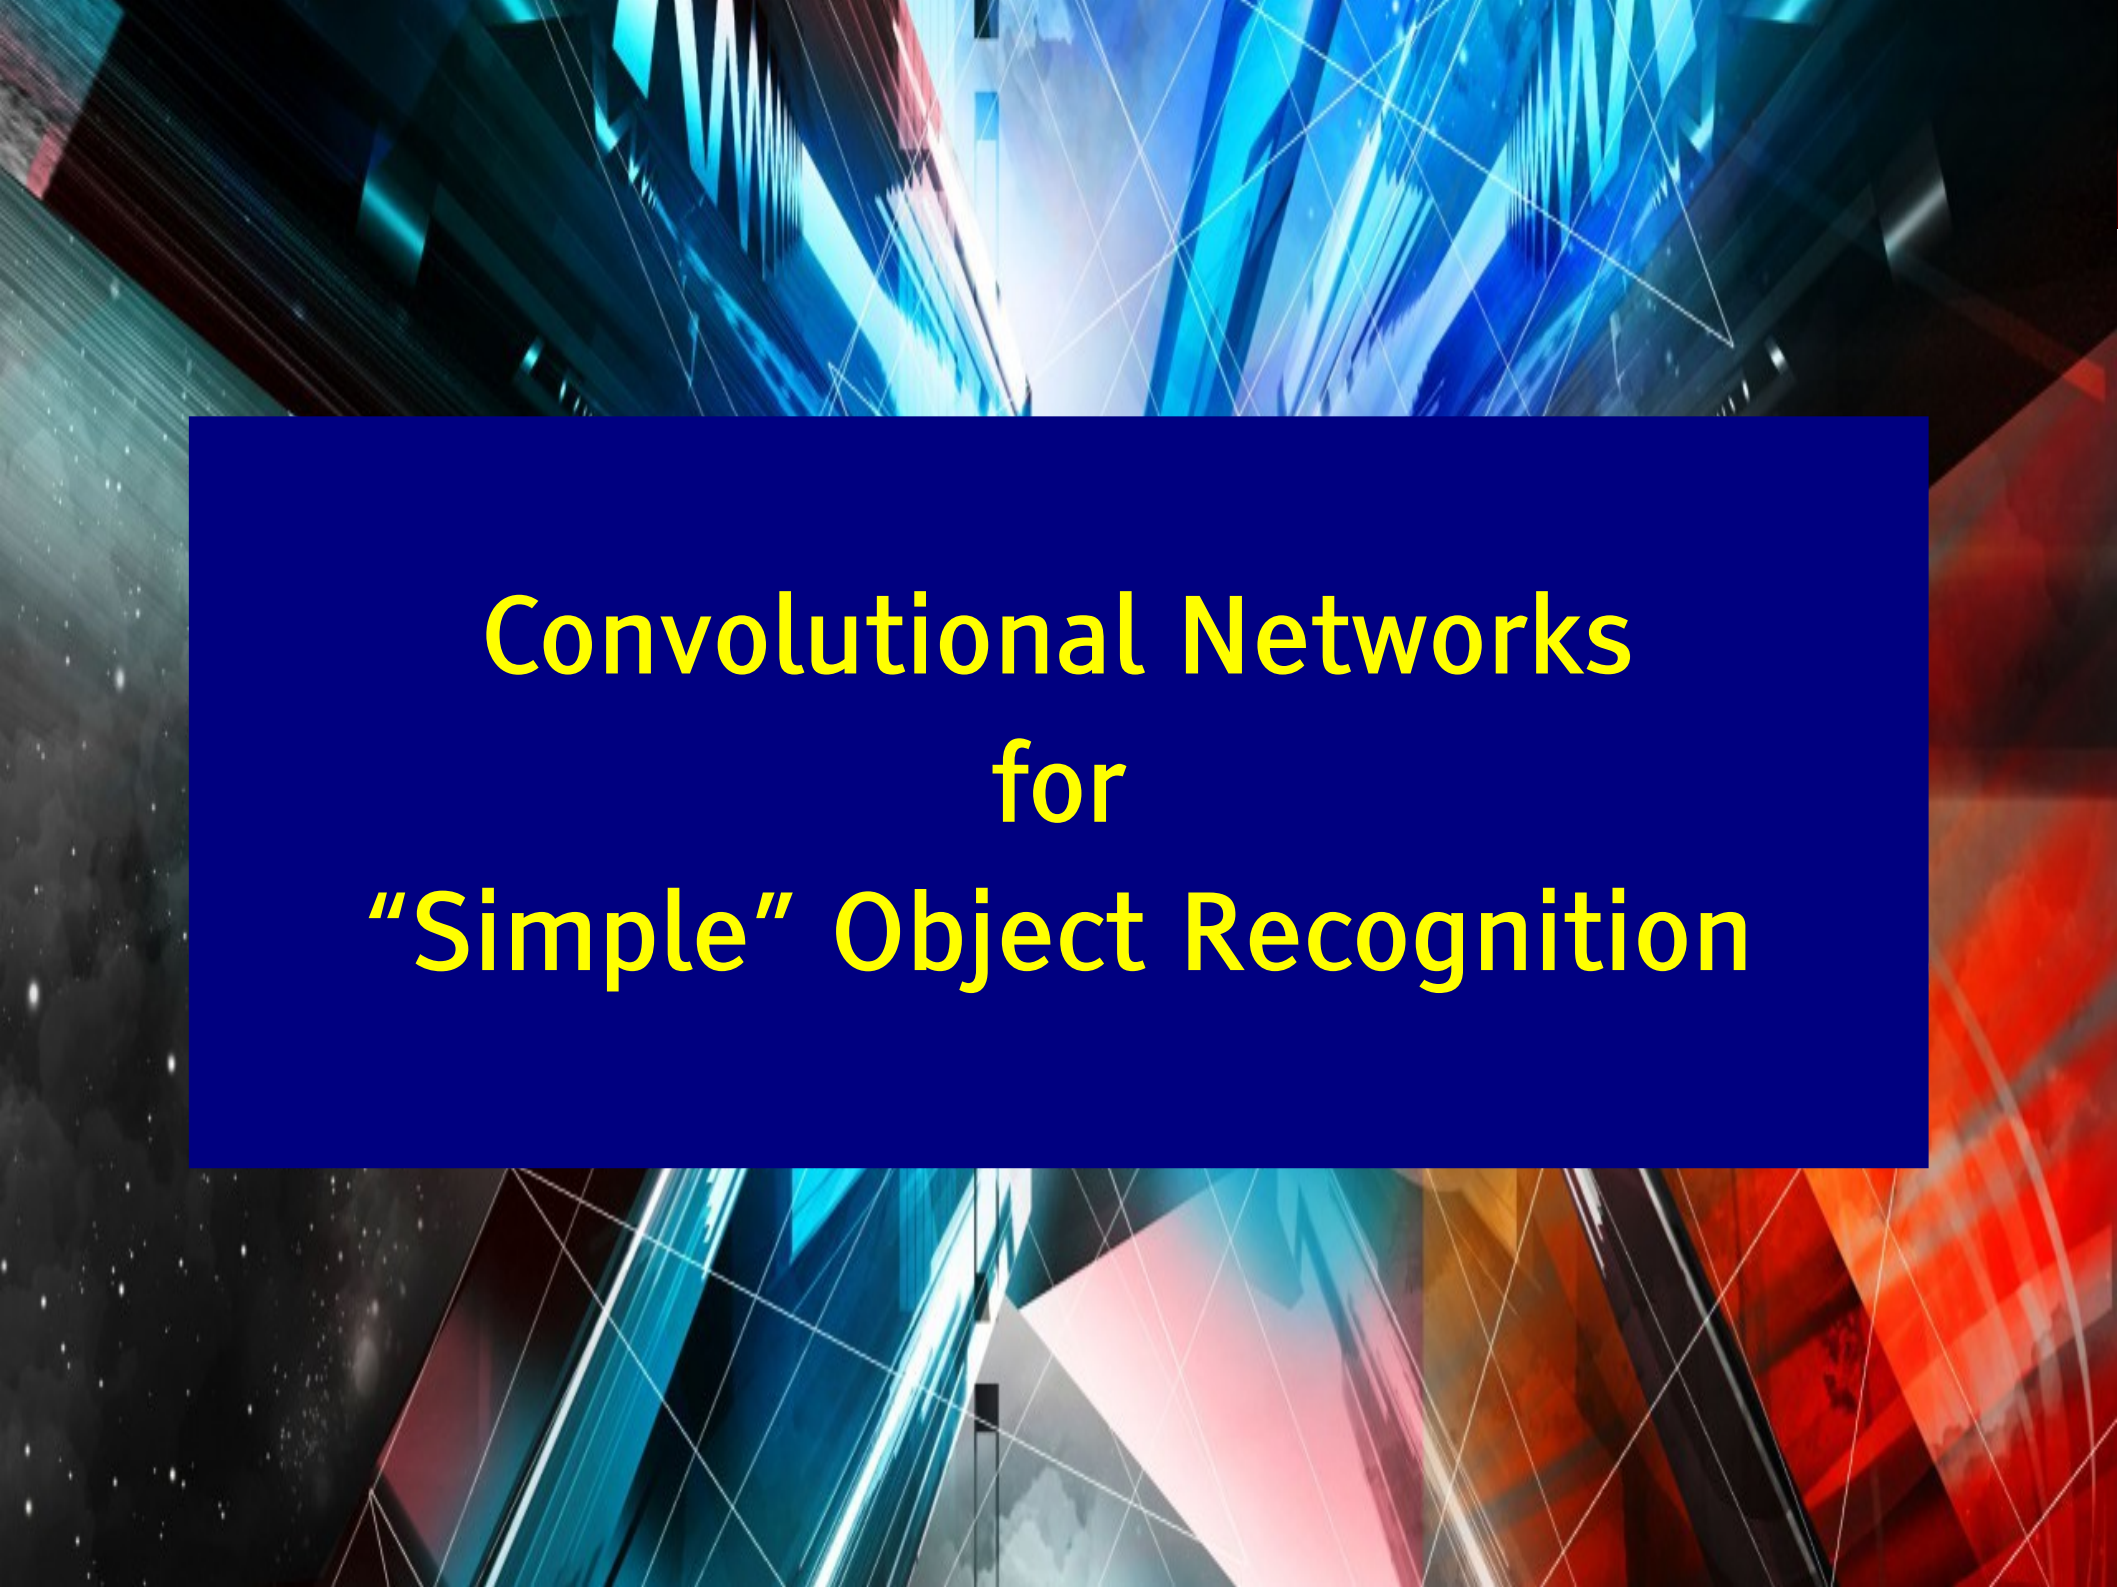
Convolutional Networks
for
“Simple” Object Recognition

We knew ConvNet worked well with characters and small images

- Traffic Sign Recognition (GTSRB)
  - German Traffic Sign Reco Bench
  - 99.2% accuracy (IDSIA)

- House Number Recognition (Google)
  - Street View House Numbers
  - 94.3 % accuracy (NYU)

NORB Dataset (2004): 5 categories, multiple views and illuminations

Y LeCun

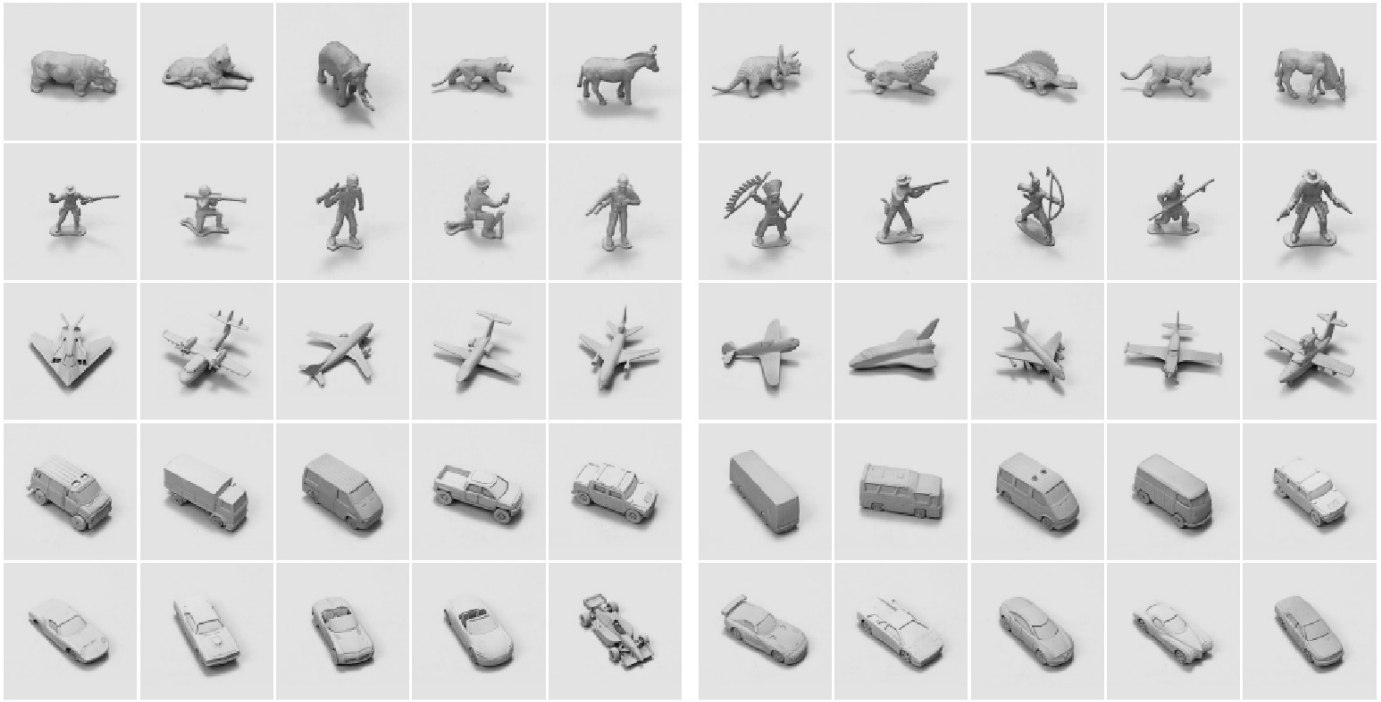

- Less than 6% error on test set with cluttered backgrounds

291,600 training samples,

58,320 test samples

- In the mid 2000s, ConvNets were getting decent results on object classification
- Dataset: "Caltech101":
  - 101 categories
  - 30 training samples per category
- But the results were slightly worse than more "traditional" computer vision methods, because
  - 1. the datasets were too small
  - 2. the computers were too slow

face

beaver

wild cat
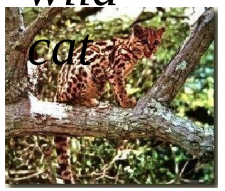

lotus

an

dollar

metronome
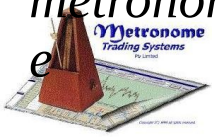

w. chair

minare

cellphon

joshua t.

cougar body

backgroun d

Late 2000s: we could get decent results on object recognition

- But we couldn't beat the state of the art because the datasets were too small
- Caltech101: 101 categories, 30 samples per category.
- But we learned that rectification and max pooling are useful! [Jarrett et al. ICCV 2009]

Single Stage System: $[64.F_{CSG}^{9\times9} - R/N/P^{5\times5}]$ **- log_reg**					
R/N/P	$R_{abs} - N - P_A$	$R_{abs} - P_A$	$N - P_M$	$N - P_A$	P_A
U^+	54.2%	50.0%	44.3%	18.5%	14.5%
R^+	54.8%	47.0%	38.0%	16.3%	14.3%
U	52.2%	43.3%(±1.6)	44.0%	17.2%	13.4%
R	53.3%	31.7%	32.1%	15.3%	12.1%(±2.2)
G	52.3%				
Two Stage System: $[64.F_{CSG}^{9\times9} - R/N/P^{5\times5}] - [256.F_{CSG}^{9\times9} - R/N/P^{4\times4}]$ **- log_reg**					
R/N/P	$R_{abs} - N - P_A$	$R_{abs} - P_A$	$N - P_M$	$N - P_A$	P_A
U^+U^+	65.5%	60.5%	61.0%	34.0%	32.0%
R^+R^+	64.7%	59.5%	60.0%	31.0%	29.7%
UU	63.7%	46.7%	56.0%	23.1%	9.1%
RR	62.9%	33.7%(±1.5)	37.6%(±1.9)	19.6%	8.8%
GT	55.8%	← like HMAX model			
Single Stage: $[64.F_{CSG}^{9\times9} - R/N/P^{5\times5}]$ **- PMK-SVM**					
U	64.0%				
Two Stages: $[64.F_{CSG}^{9\times9} - R/N/P^{5\times5}] - [256.F_{CSG}^{9\times9} - R/N]$ **- PMK-SVM**					
UU	52.8%				

ConvNet for Segmentation
Self-Learning ConvNet for
Vision-Based Navigation
for Off-Road Robots

ConvNets for Image Segmentation

- Biological Image Segmentation
  - [Ning et al. IEEE-TIP 2005]
- Pixel labeling with large context using a convnet
- ConvNet takes a window of pixels and produces a label for the central pixel
- Cleanup using a kind of conditional random field (CRF)
  - Similar to a field of expert

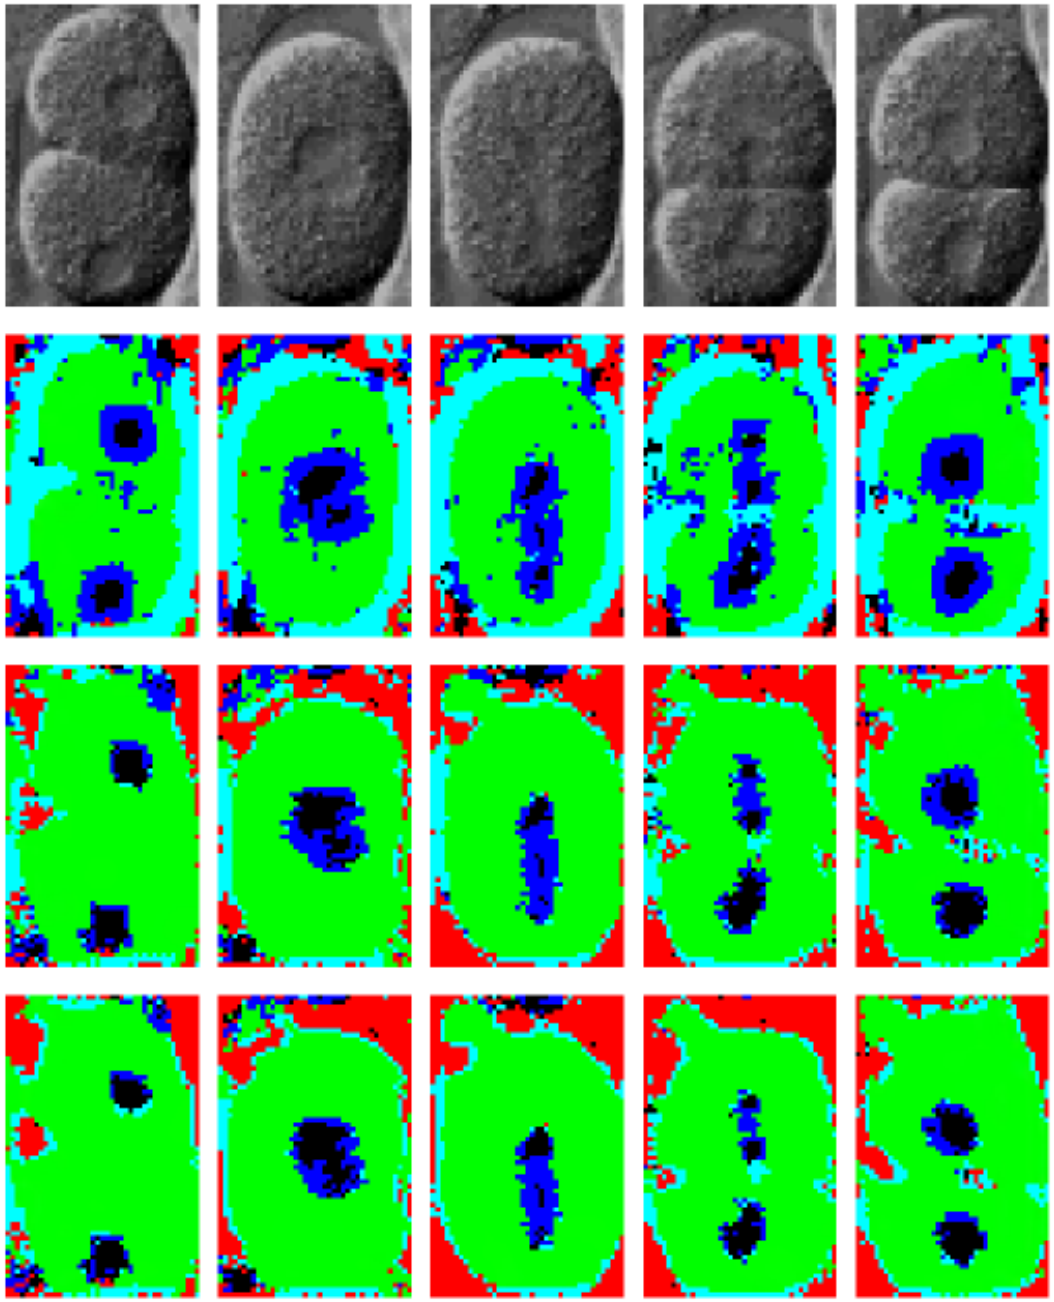

ConvNet in Connectomics [Jain, Turaga, Seung 2007-present]

- 3D ConvNet Volumetric Images
- Each voxel labeled as "membrane" or "non-membrane using a 7x7x7 voxel neighborhood
- Has become a standard method in connectomics

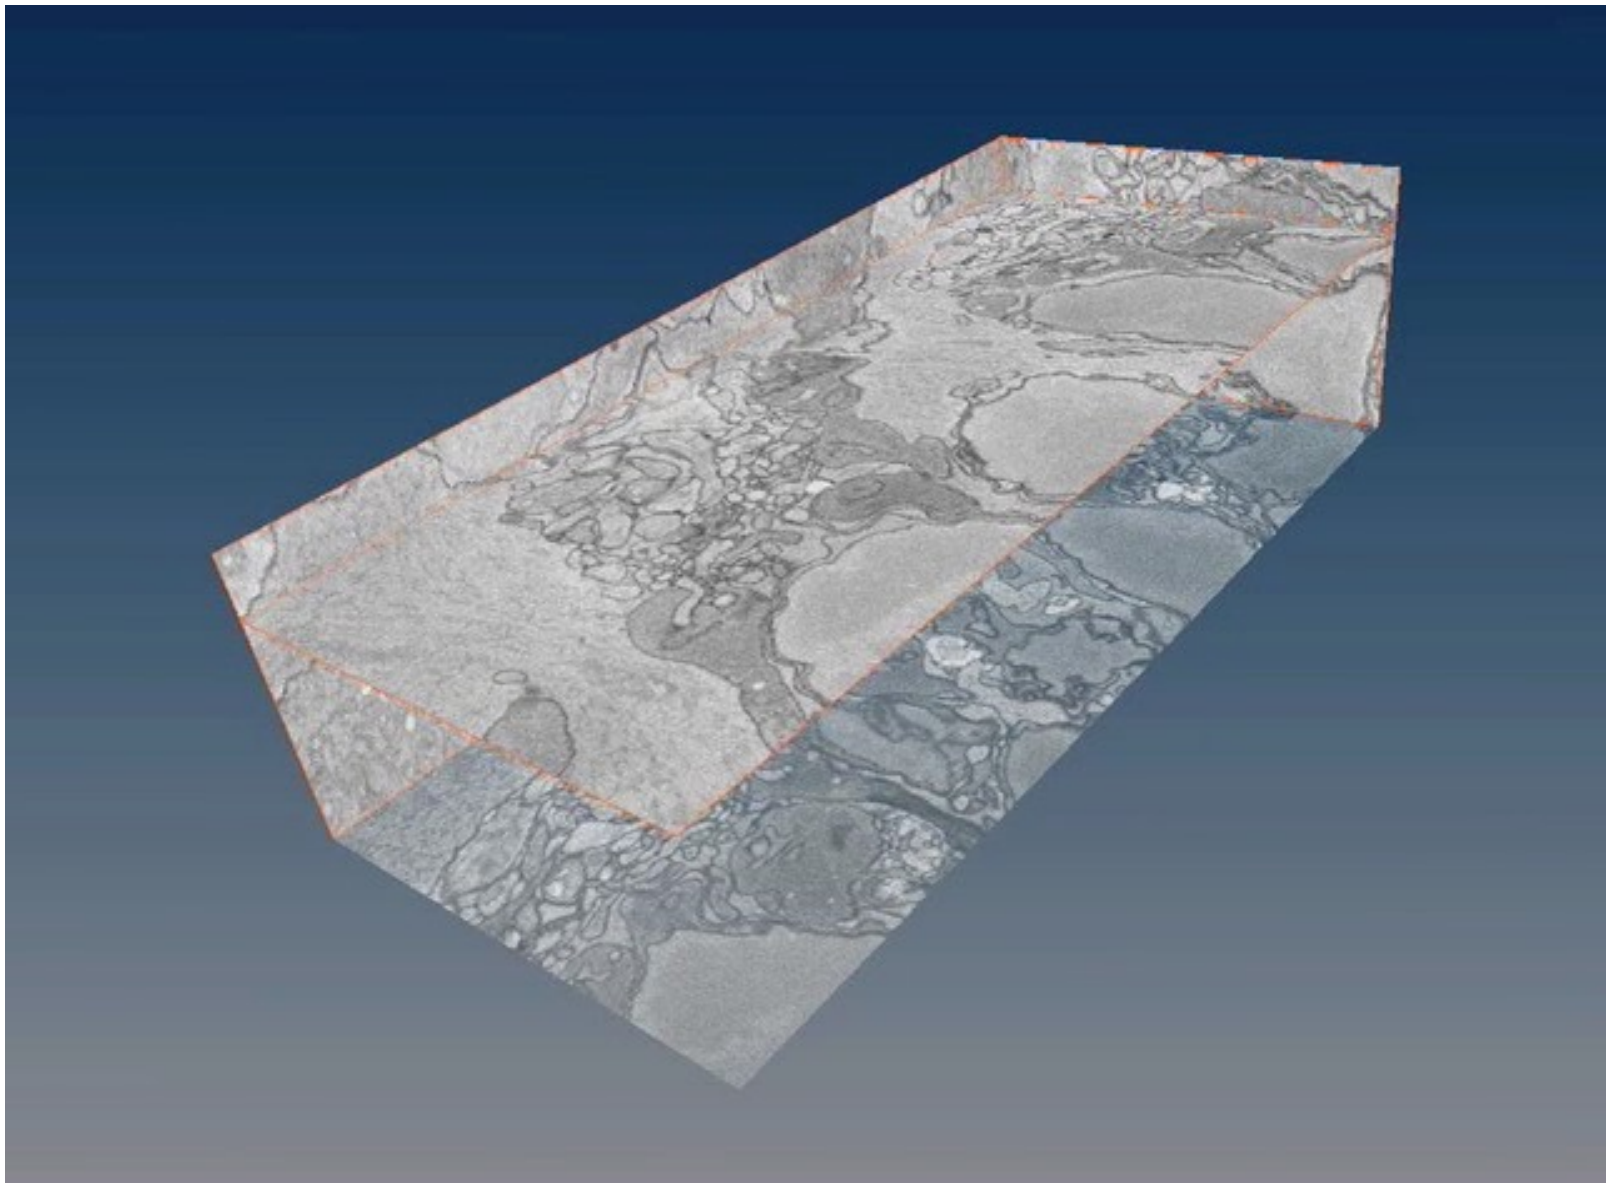

LAGR project (2005-2008): vision-based navigation for off-road robot

Y LeCun

- Getting a robot to drive autonomously in unknown terrain solely from vision (camera input).
- Our team (NYU/Net-Scale Technologies Inc.) was one of 8 participants funded by DARPA
- All teams received identical robots and can only modify the software (not the hardware)
- The robot is given the GPS coordinates of a goal, and must drive to the goal as fast as possible. The terrain is unknown in advance. The robot is run 3 times through the same course.
- Long-Range Obstacle Detection with on-line, self-trained ConvNet
- Uses temporal consistency!

Obstacle Detection at Short Range: Stereovision

Obstacles overlaid with camera image

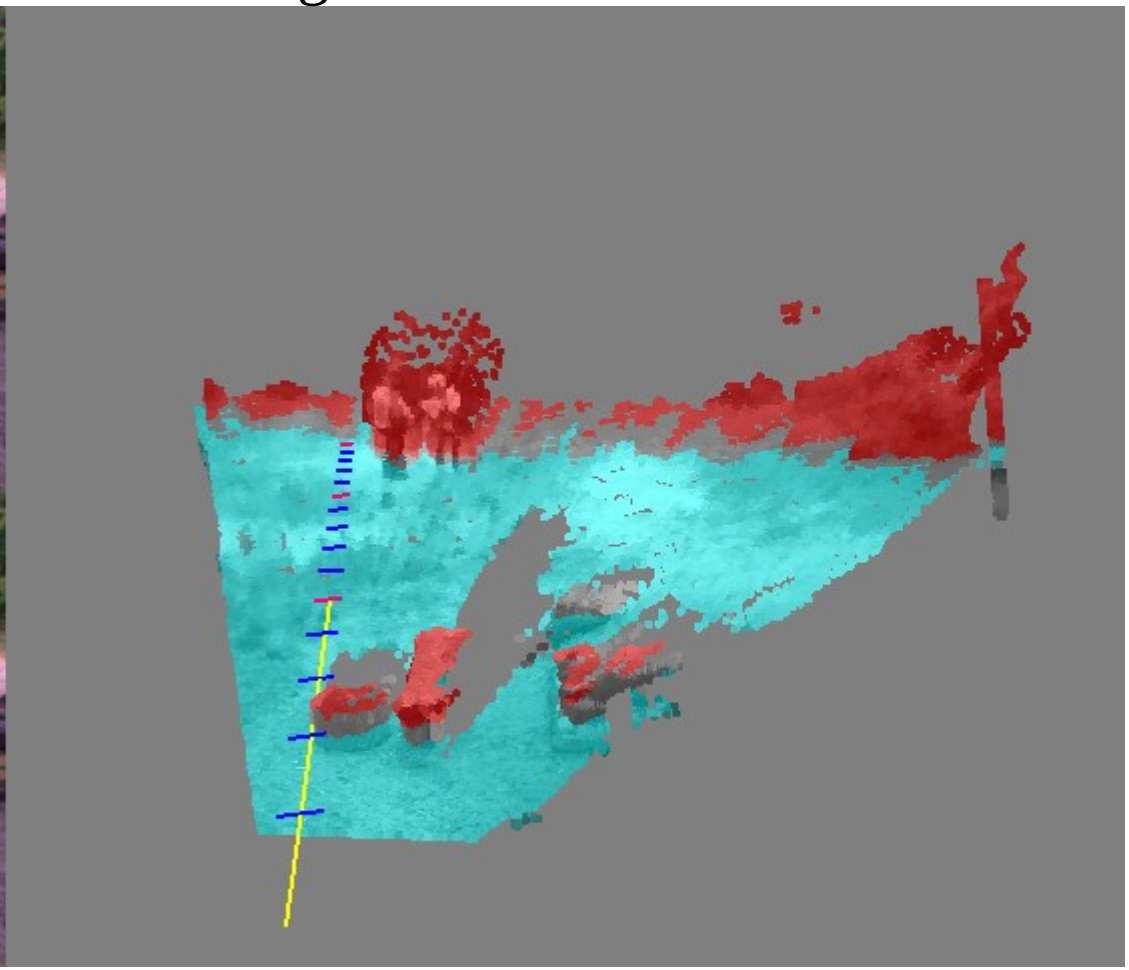

Camera image

Detected obstacles (red)

But Stereovision Doesn't work at long range

- Stereo is only good up to about 10 meters.
- But not seeing past 10 meters is like driving in a fog or a snowstorm!

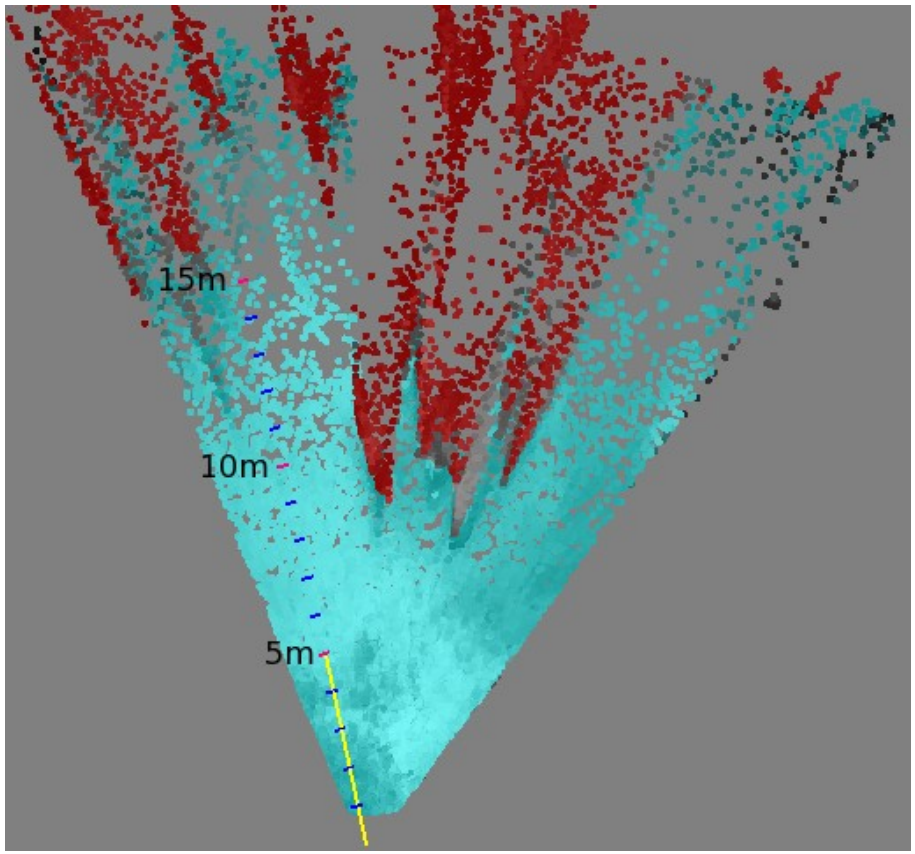

Long Range Vision with a Convolutional Net

Y LeCun

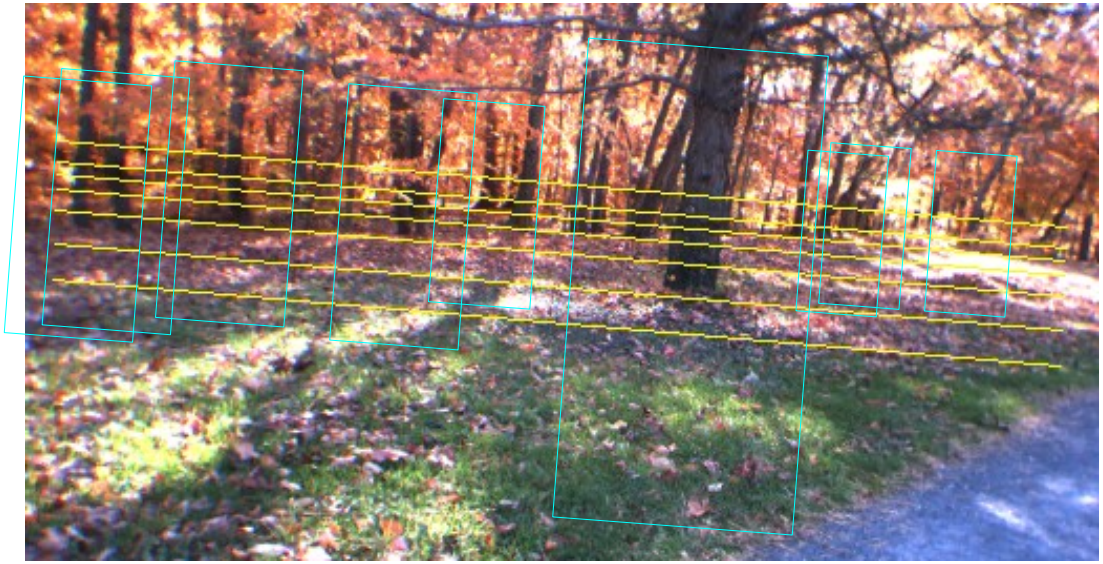

Pre-processing (125 ms)

- Ground plane estimation
- Horizon leveling
- Conversion to YUV + local contrast normalization
- Scale invariant pyramid of distance-normalized image "bands"

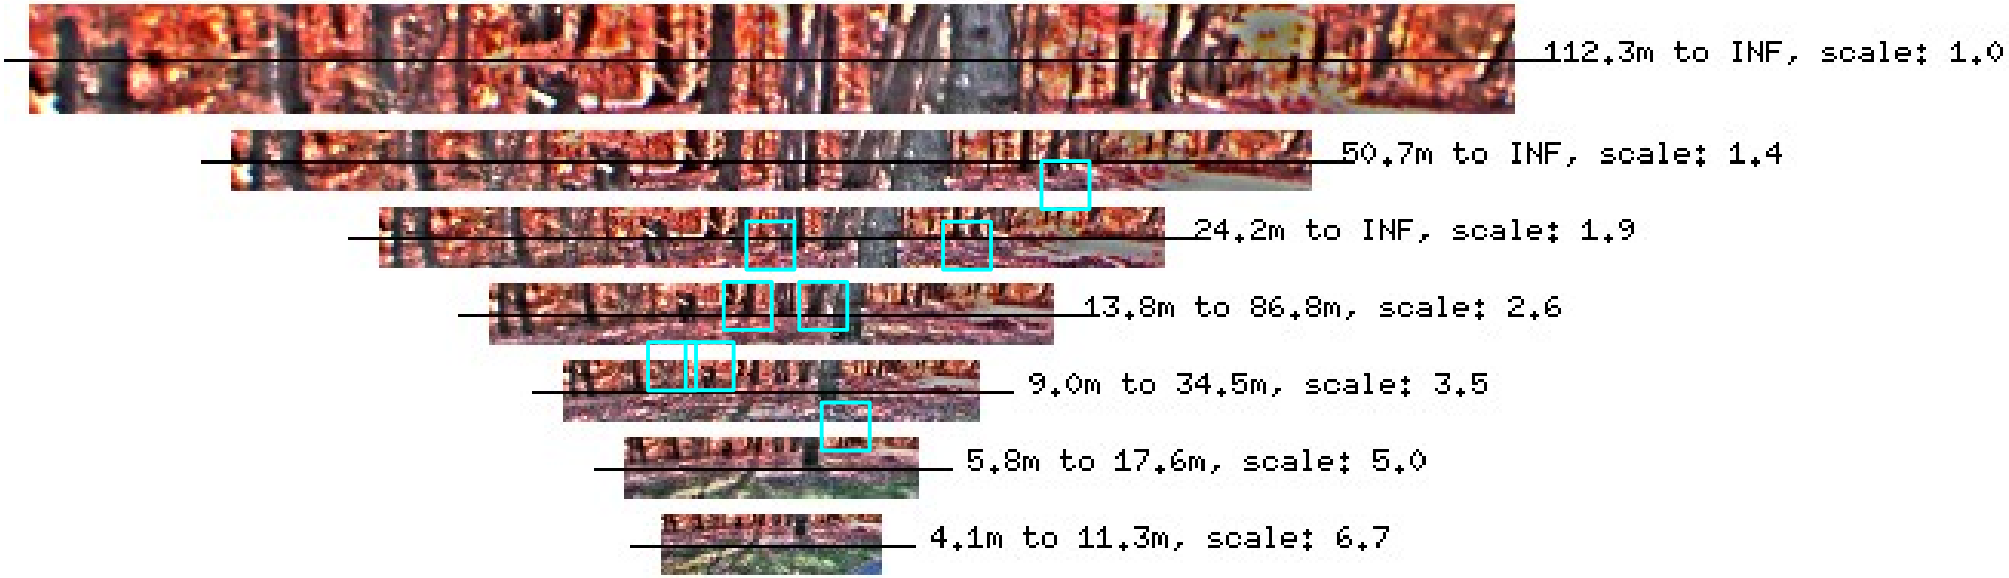

Convolutional Net Architecture

100 features per
3x12x25 input window

100@25x121

CONVOLUTIONS (6x5)

YUV image band
20-36 pixels tall,
36-500 pixels wide

20@30x125

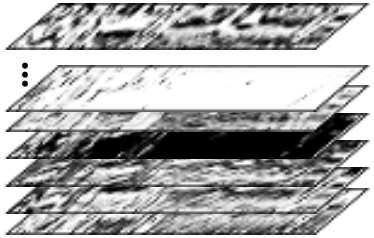

MAX SUBSAMPLING (1x4)

20@30x484

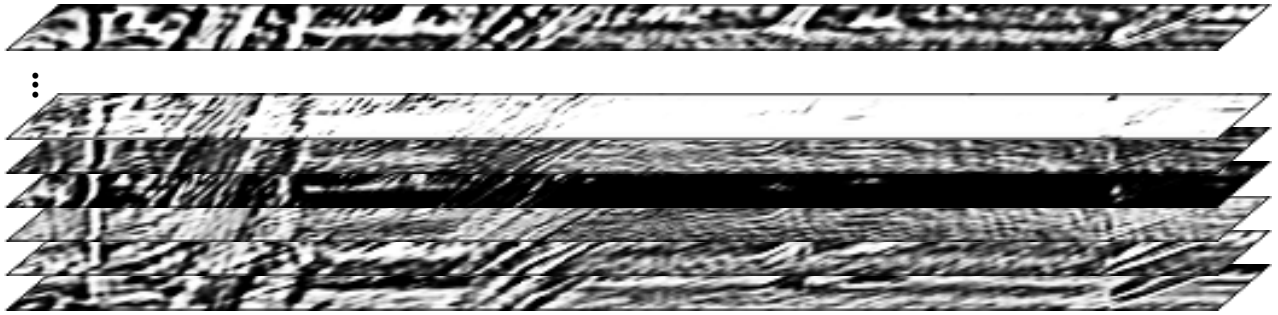

CONVOLUTIONS (7x6)

3@36x484

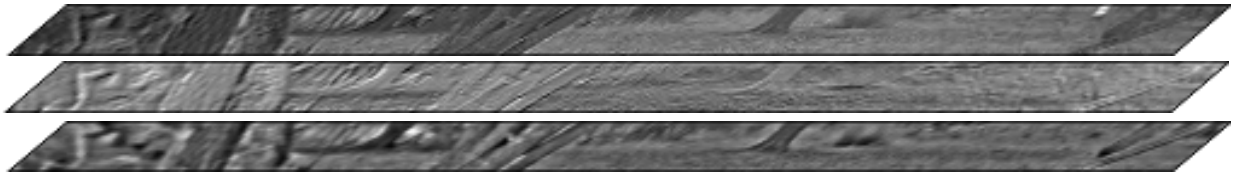

YUV input

Scene Labeling with ConvNet + online learning

- Image Labeling for Off-Road Robots [Hadsell JFR 2008]
  - ConvNet labels pixels as one of 3 categories
  - Traversible/flat (green), non traversible (red), foot of obstacle (purple)
  - Labels obtained from stereo vision and SLAM

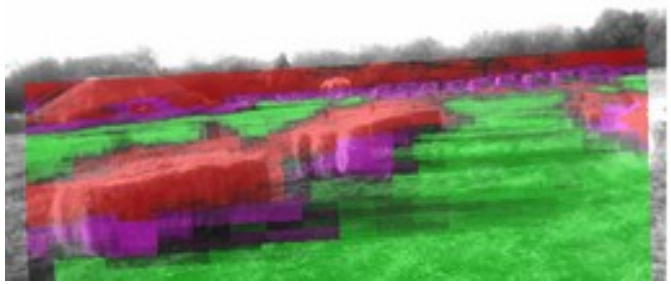

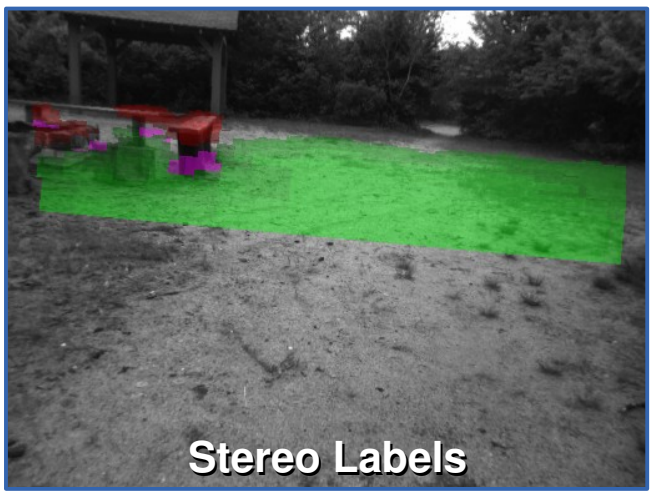

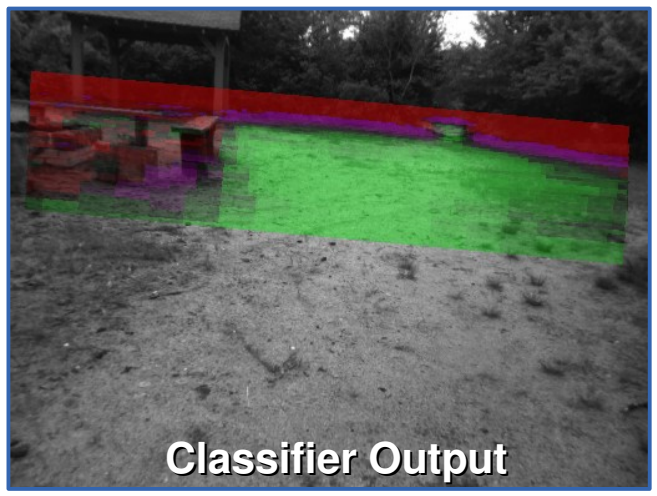

Long Range Vision Results

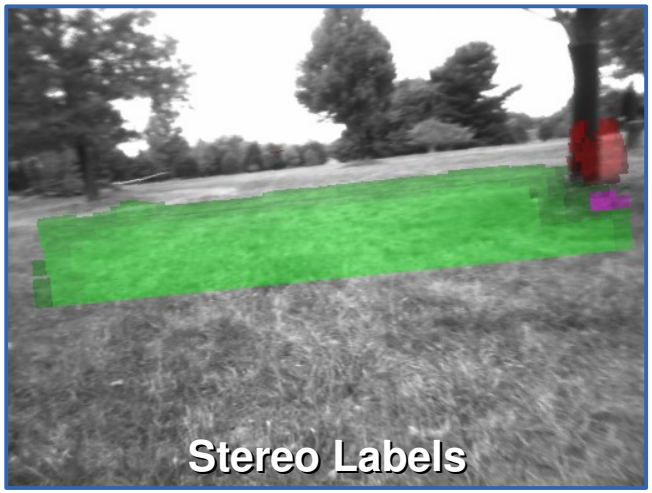

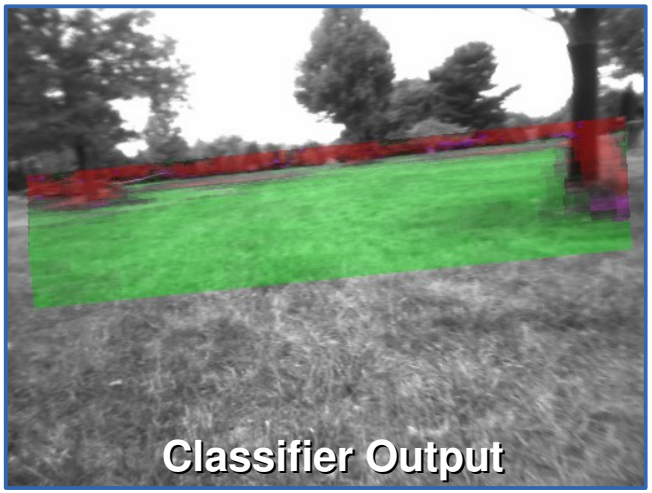

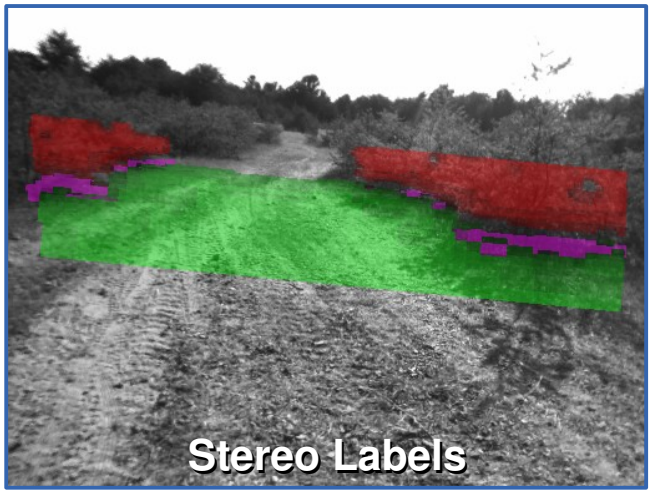

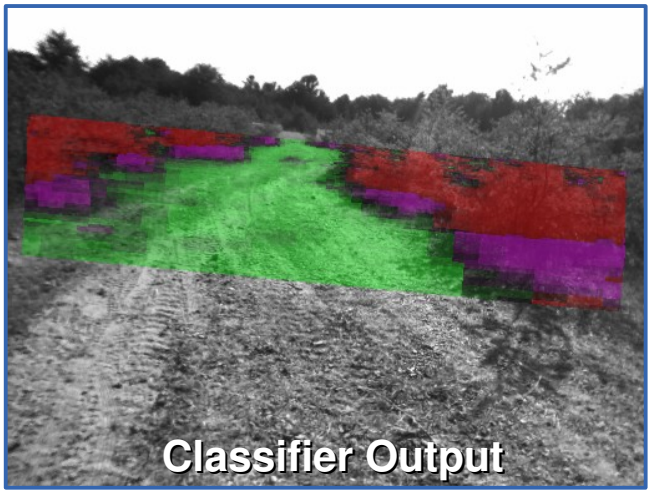

Long Range Vision Results

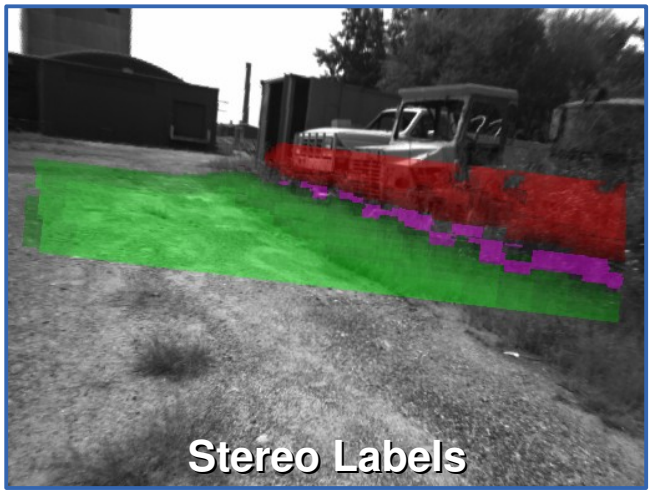

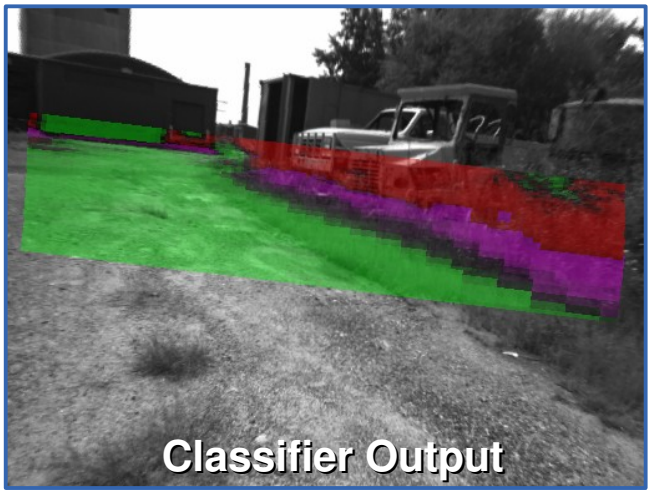

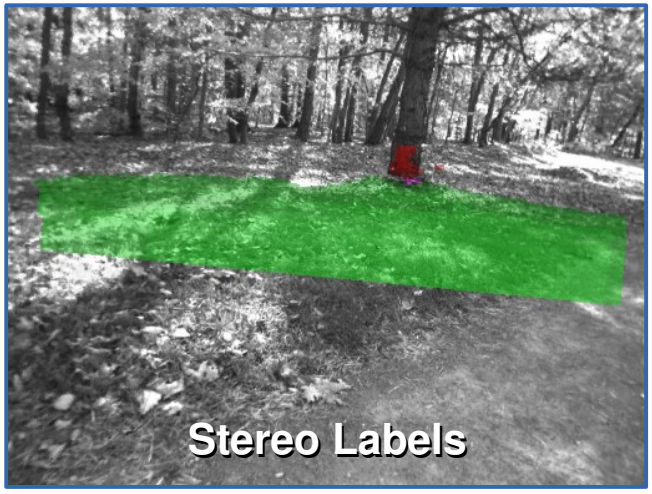

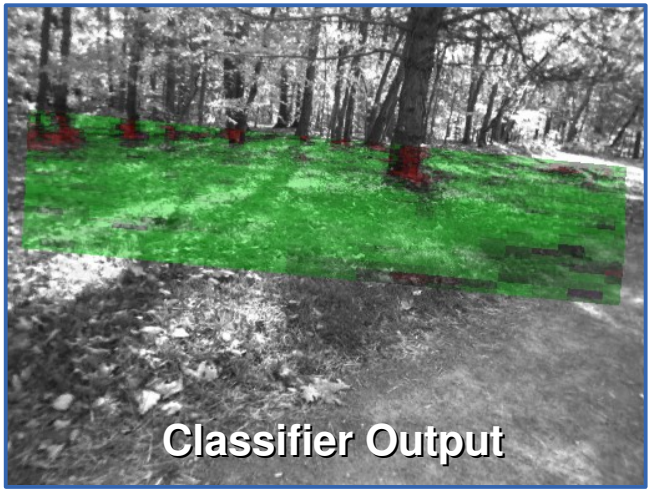

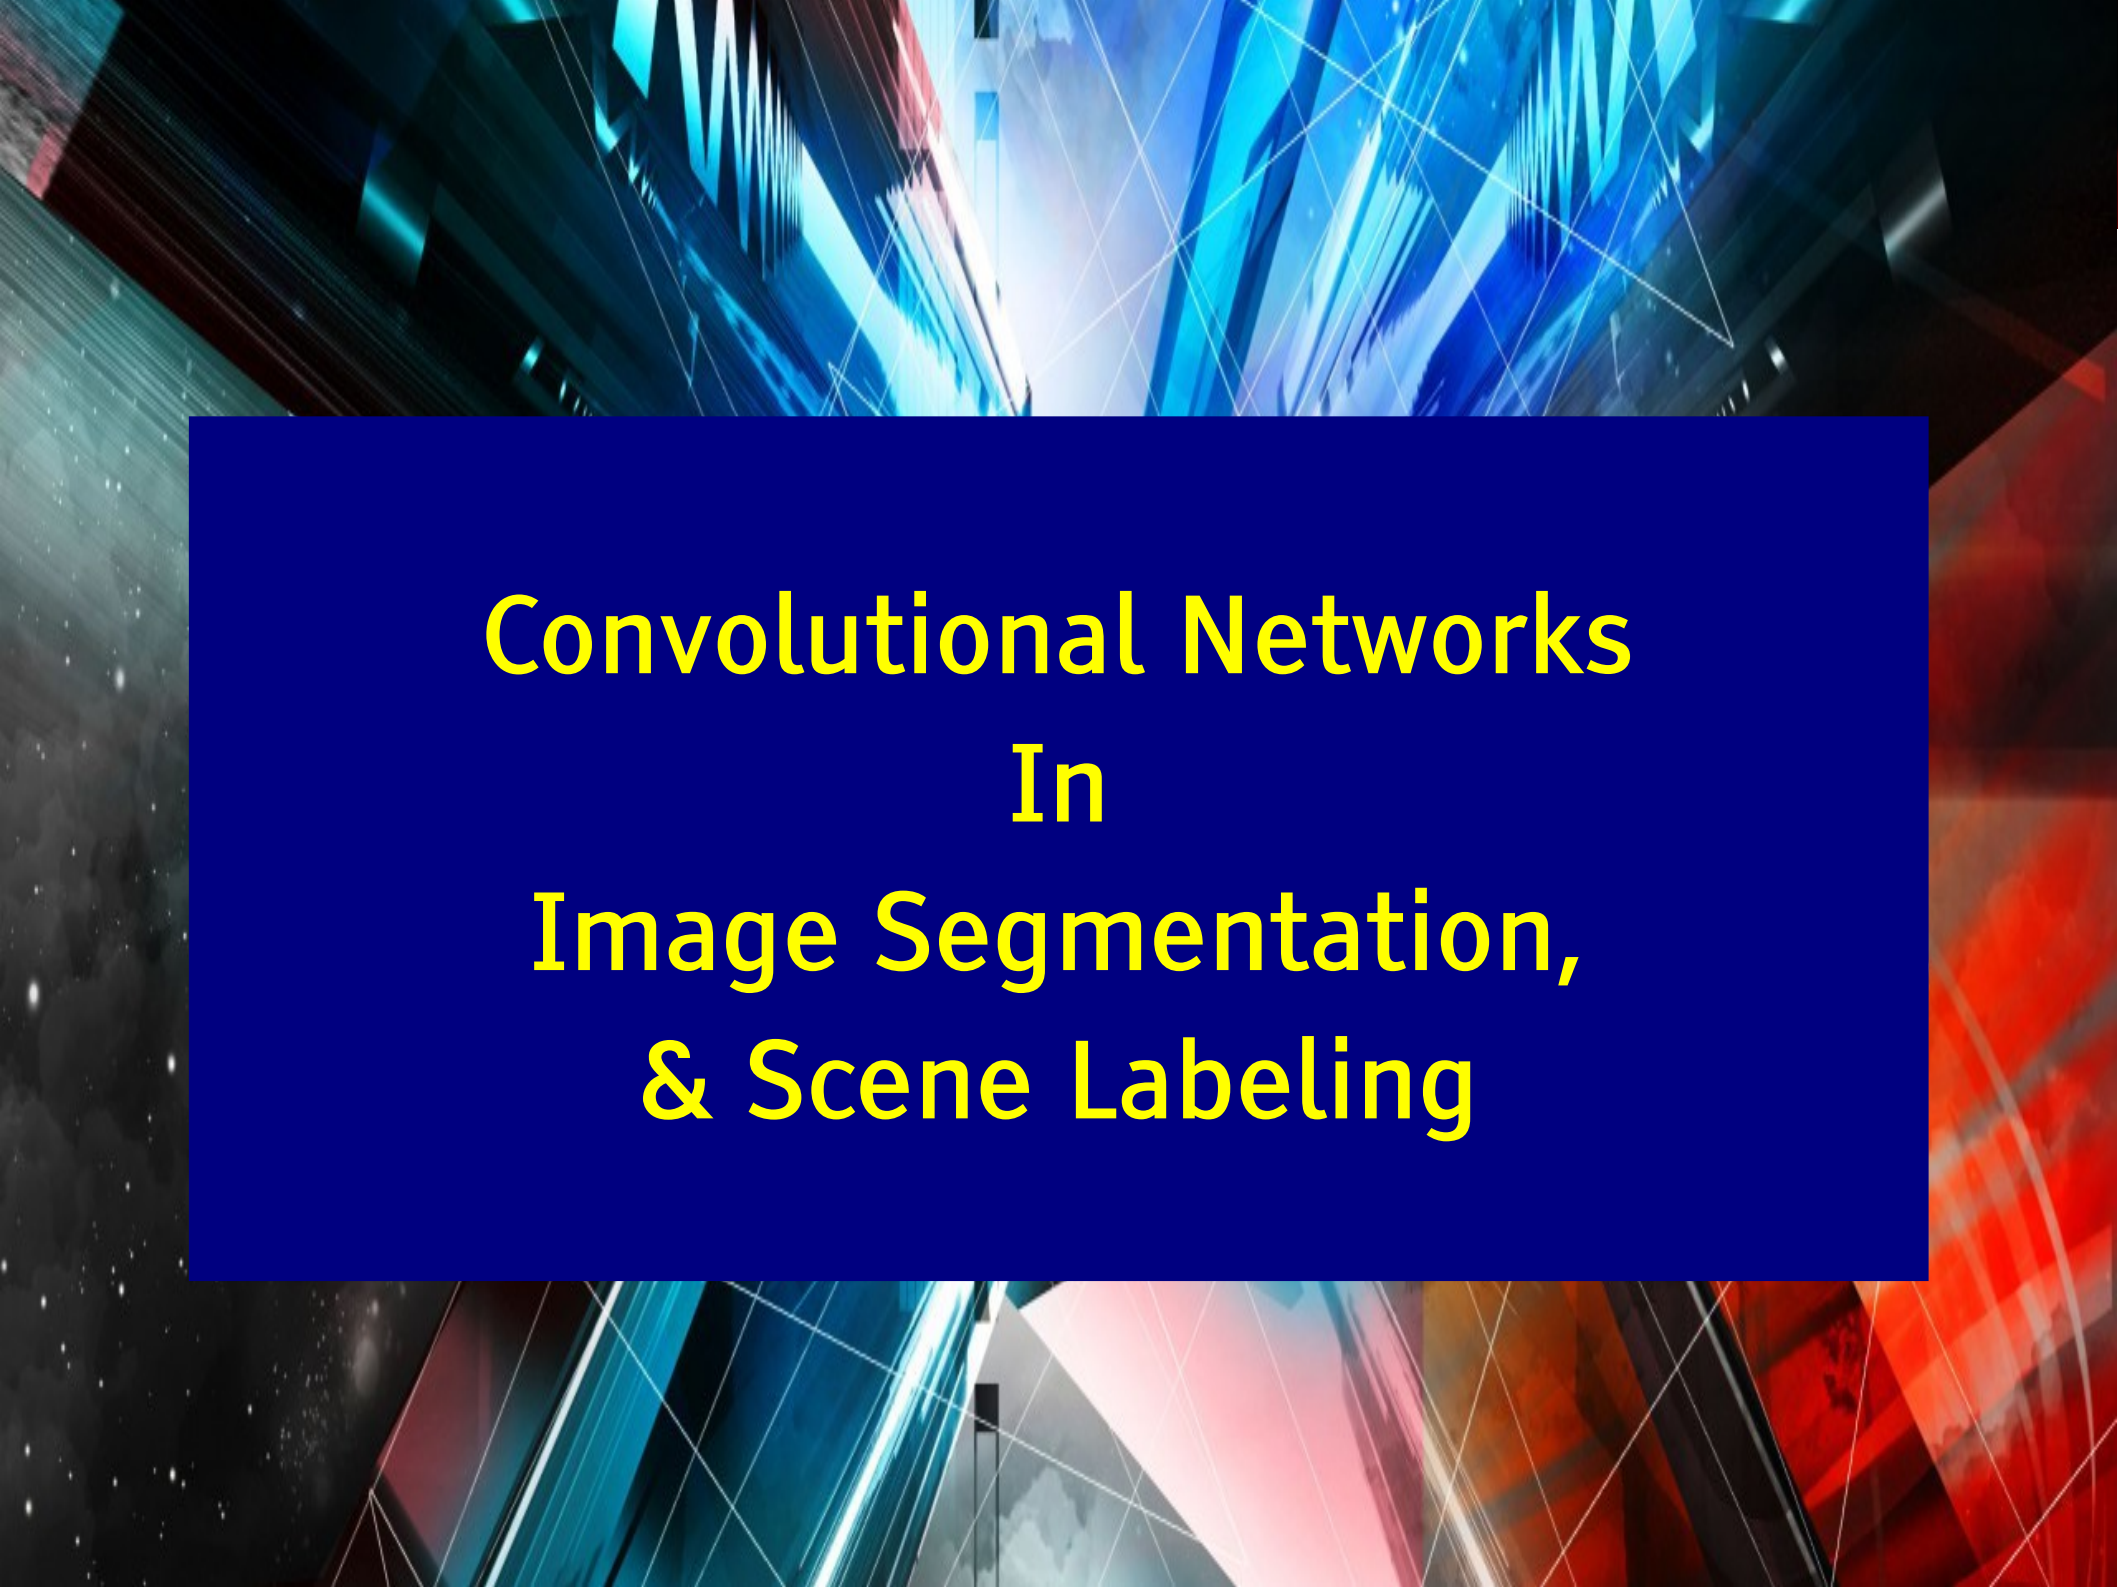
Convolutional Networks
In
Image Segmentation,
& Scene Labeling

Semantic Labeling / Scene Parsing: Labeling every pixel with the object it belongs to

- Would help identify obstacles, targets, landing sites, dangerous areas
- Would help line up depth map with edge maps

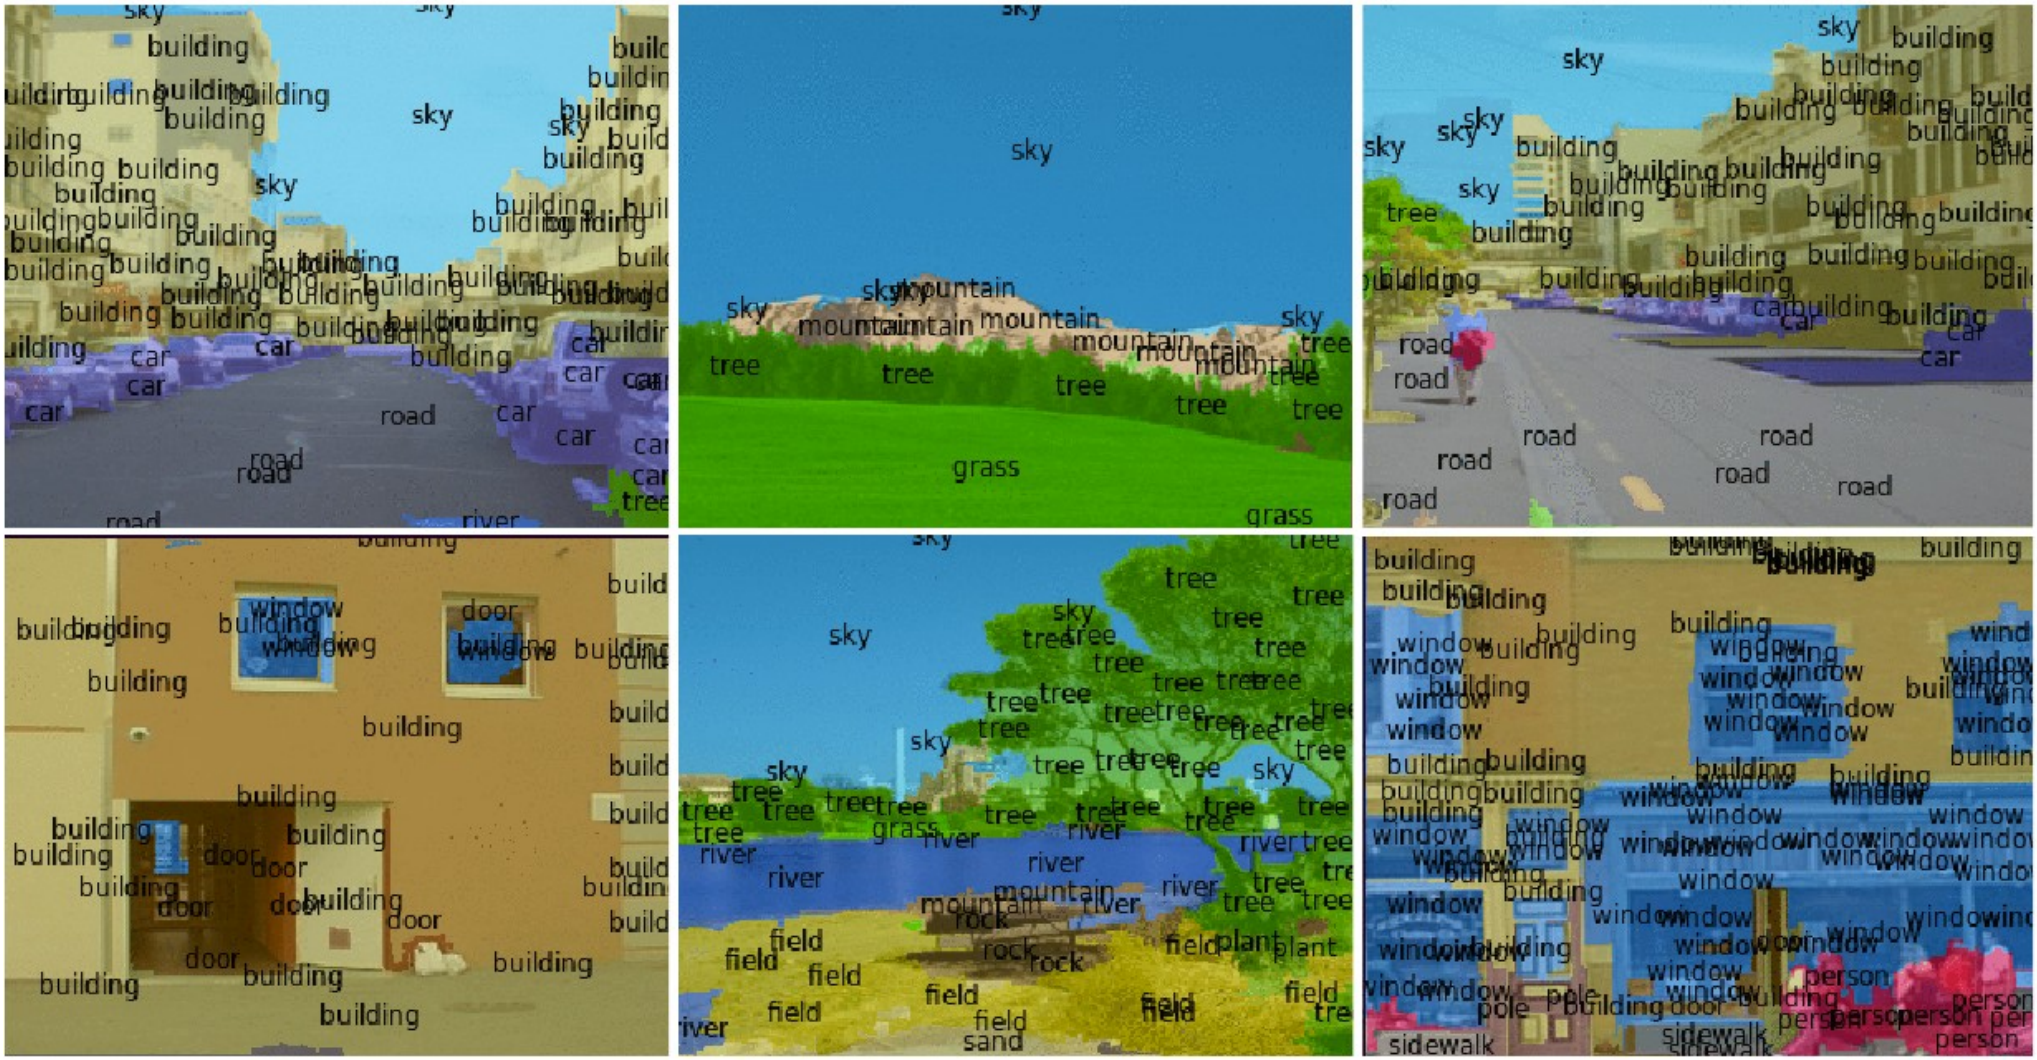

[Farabet et al. ICML 2012, PAMI 2013]

Scene Parsing/Labeling: ConvNet Architecture

- Each output sees a large input context:
  - **46x46** window at full rez; **92x92** at ½ rez; **184x184** at ¼ rez
  - [7x7conv]->[2x2pool]->[7x7conv]->[2x2pool]->[7x7conv]->
  - Trained supervised on fully-labeled images

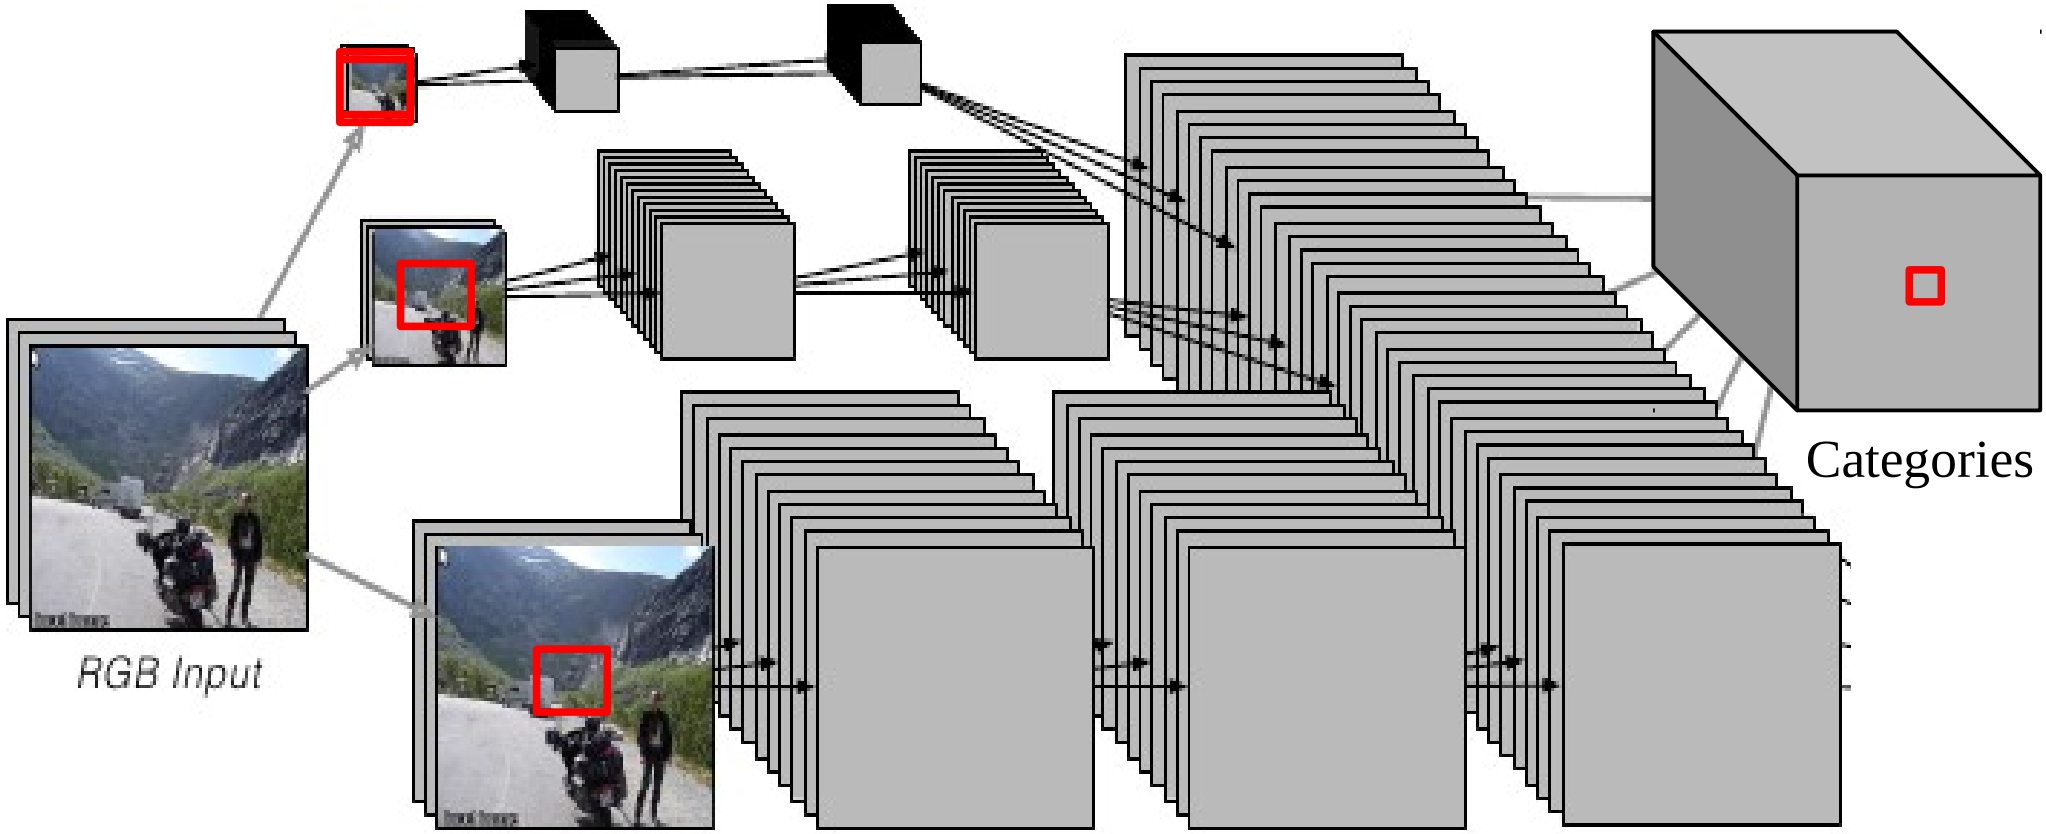

Method 1: majority over super-pixel regions

Super-pixel boundary hypetheses

Superpixel boundaries

Majority Vote Over Superpixels

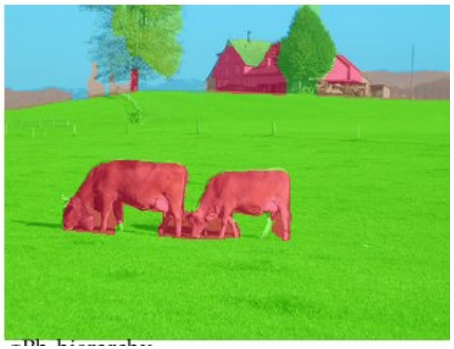

Categories aligned With region boundaries

Input image

Multi-scale ConvNet

Features from Convolutional net (d=768 per pixel)

Convolutional classifier

“soft” categories scores

[Farabet et al. IEEE T. PAMI 2013]

Scene Parsing/Labeling: Performance

- Stanford Background Dataset [Gould 1009]: 8 categories

	Pixel Acc.	Class Acc.	CT (sec.)
Gould *et al.* 2009 [14]	76.4%	-	10 to 600s
Munoz *et al.* 2010 [32]	76.9%	66.2%	12s
Tighe *et al.* 2010 [46]	77.5%	-	10 to 300s
Socher *et al.* 2011 [45]	78.1%	-	?
Kumar *et al.* 2010 [22]	79.4%	-	$< 600s$
Lempitzky *et al.* 2011 [28]	**81.9%**	72.4%	$> 60s$
singlescale convnet	66.0 %	56.5 %	0.35s
multiscale convnet	78.8 %	72.4%	0.6s
multiscale net + superpixels	80.4%	74.56%	**0.7s**
multiscale net + gPb + cover	80.4%	75.24%	61s
multiscale net + CRF on gPb	81.4%	**76.0%**	60.5s

[Farabet et al. IEEE T. PAMI 2013]

Scene Parsing/Labeling: Performance

	Pixel Acc.	Class Acc.
Liu *et al.* 2009 [31]	74.75%	-
Tighe *et al.* 2010 [44]	76.9%	29.4%
raw multiscale net[1]	67.9%	45.9%
multiscale net + superpixels[1]	71.9%	**50.8%**
multiscale net + cover[1]	72.3%	**50.8%**
multiscale net + cover[2]	**78.5%**	29.6%

- SIFT Flow Dataset
- [Liu 2009]:
- 33 categories

- Barcelona dataset
- [Tighe 2010]:
- 170 categories.

	Pixel Acc.	Class Acc.
Tighe *et al.* 2010 [44]	66.9%	7.6%
raw multiscale net[1]	37.8%	**12.1%**
multiscale net + superpixels[1]	44.1%	**12.4%**
multiscale net + cover[1]	46.4%	**12.5%**
multiscale net + cover[2]	**67.8%**	**9.5%**

[Farabet et al. IEEE T. PAMI 2012]

Scene Parsing/Labeling: SIFT Flow dataset (33 categories)

- Samples from the SIFT-Flow dataset (Liu)

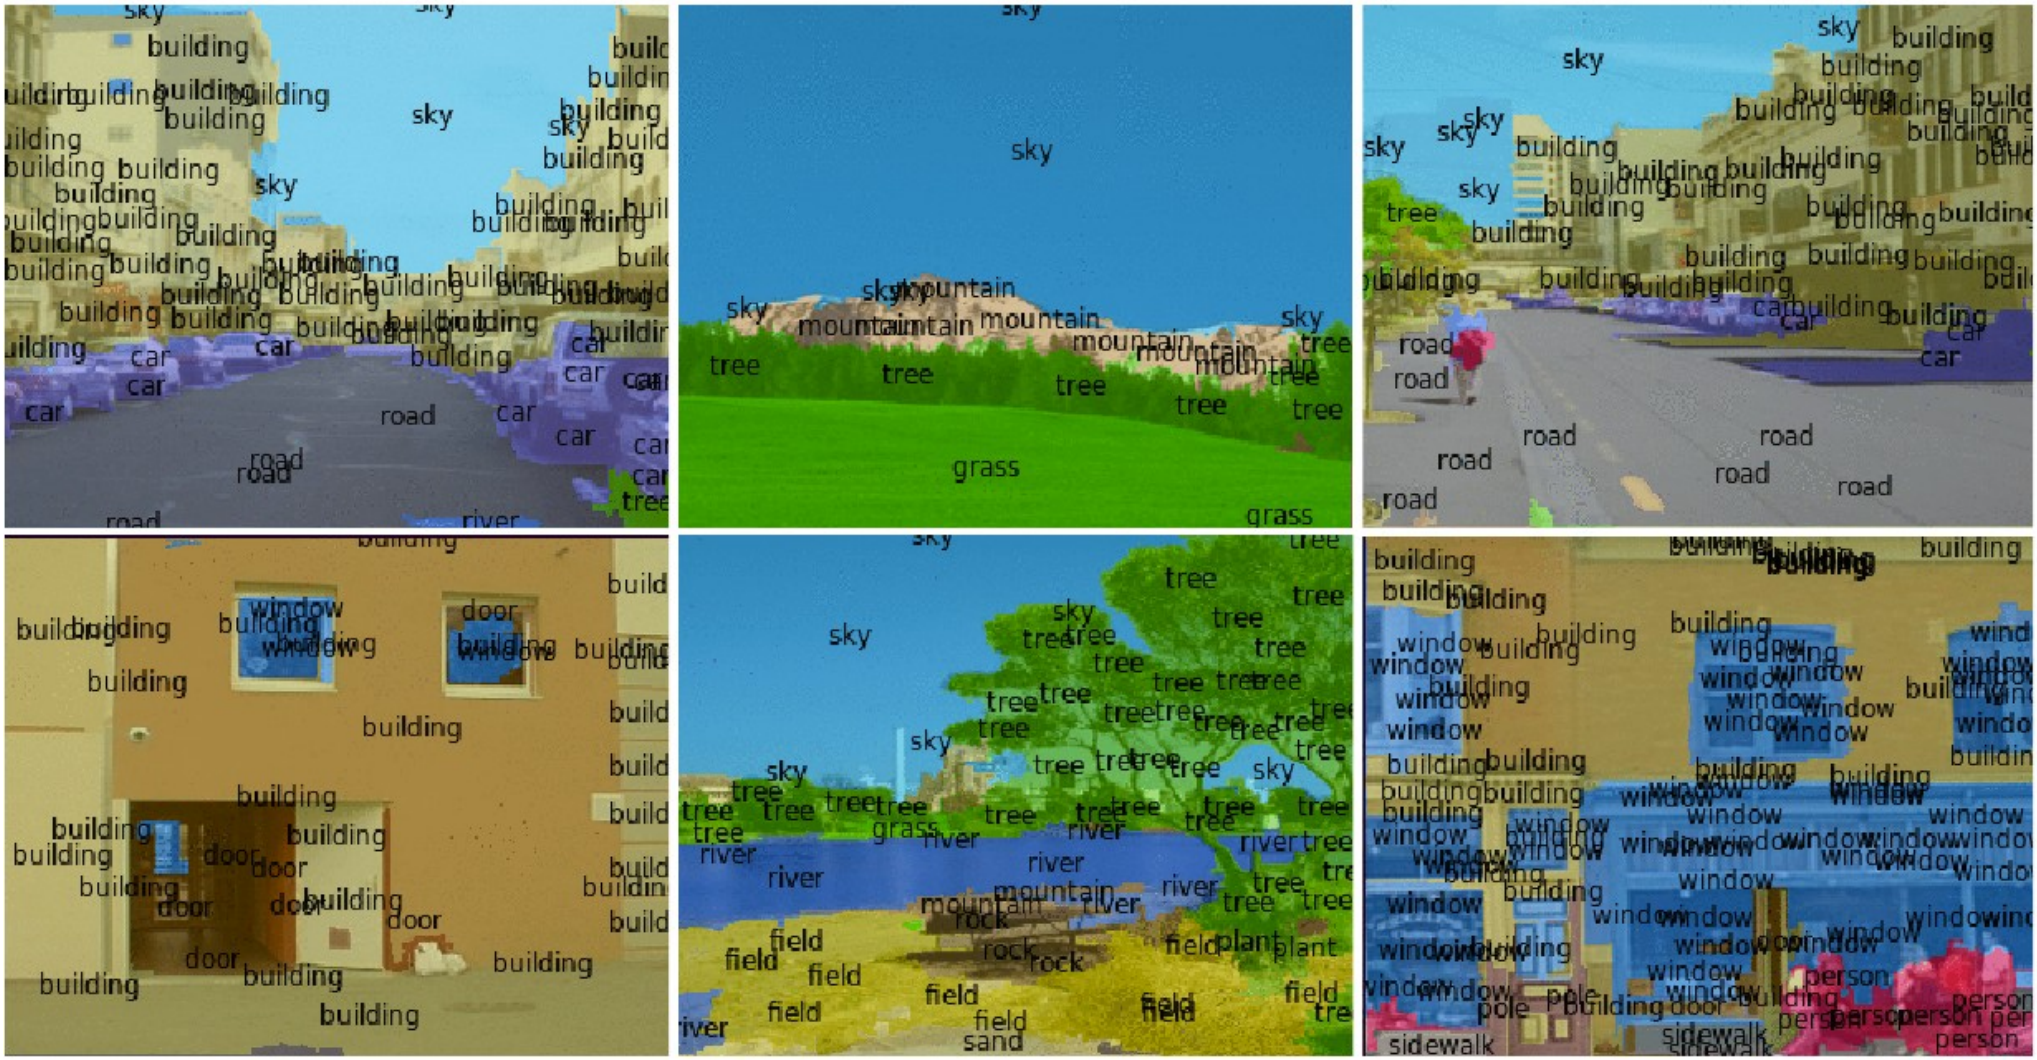

[Farabet et al. ICML 2012, PAMI 2013]

Scene Parsing/Labeling: SIFT Flow dataset (33 categories)

Y LeCun

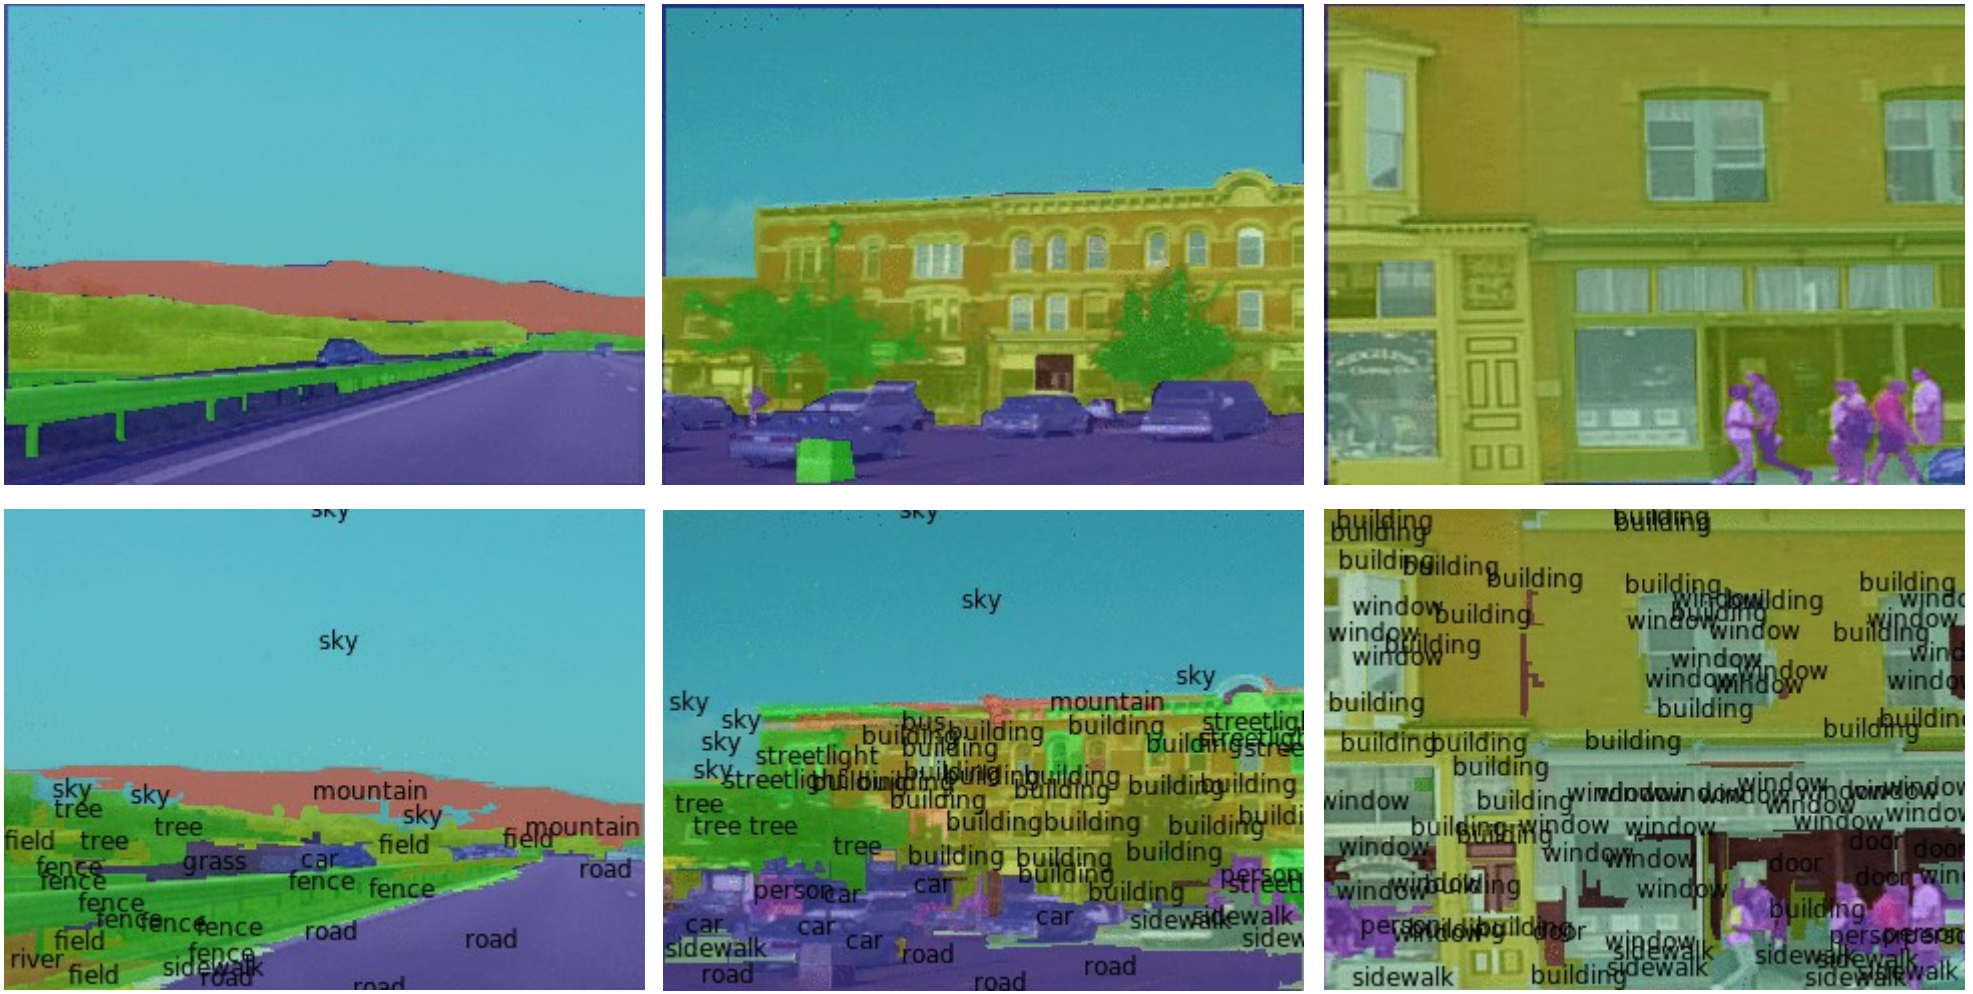

[Farabet et al. ICML 2012, PAMI 2013]

Scene Parsing/Labeling

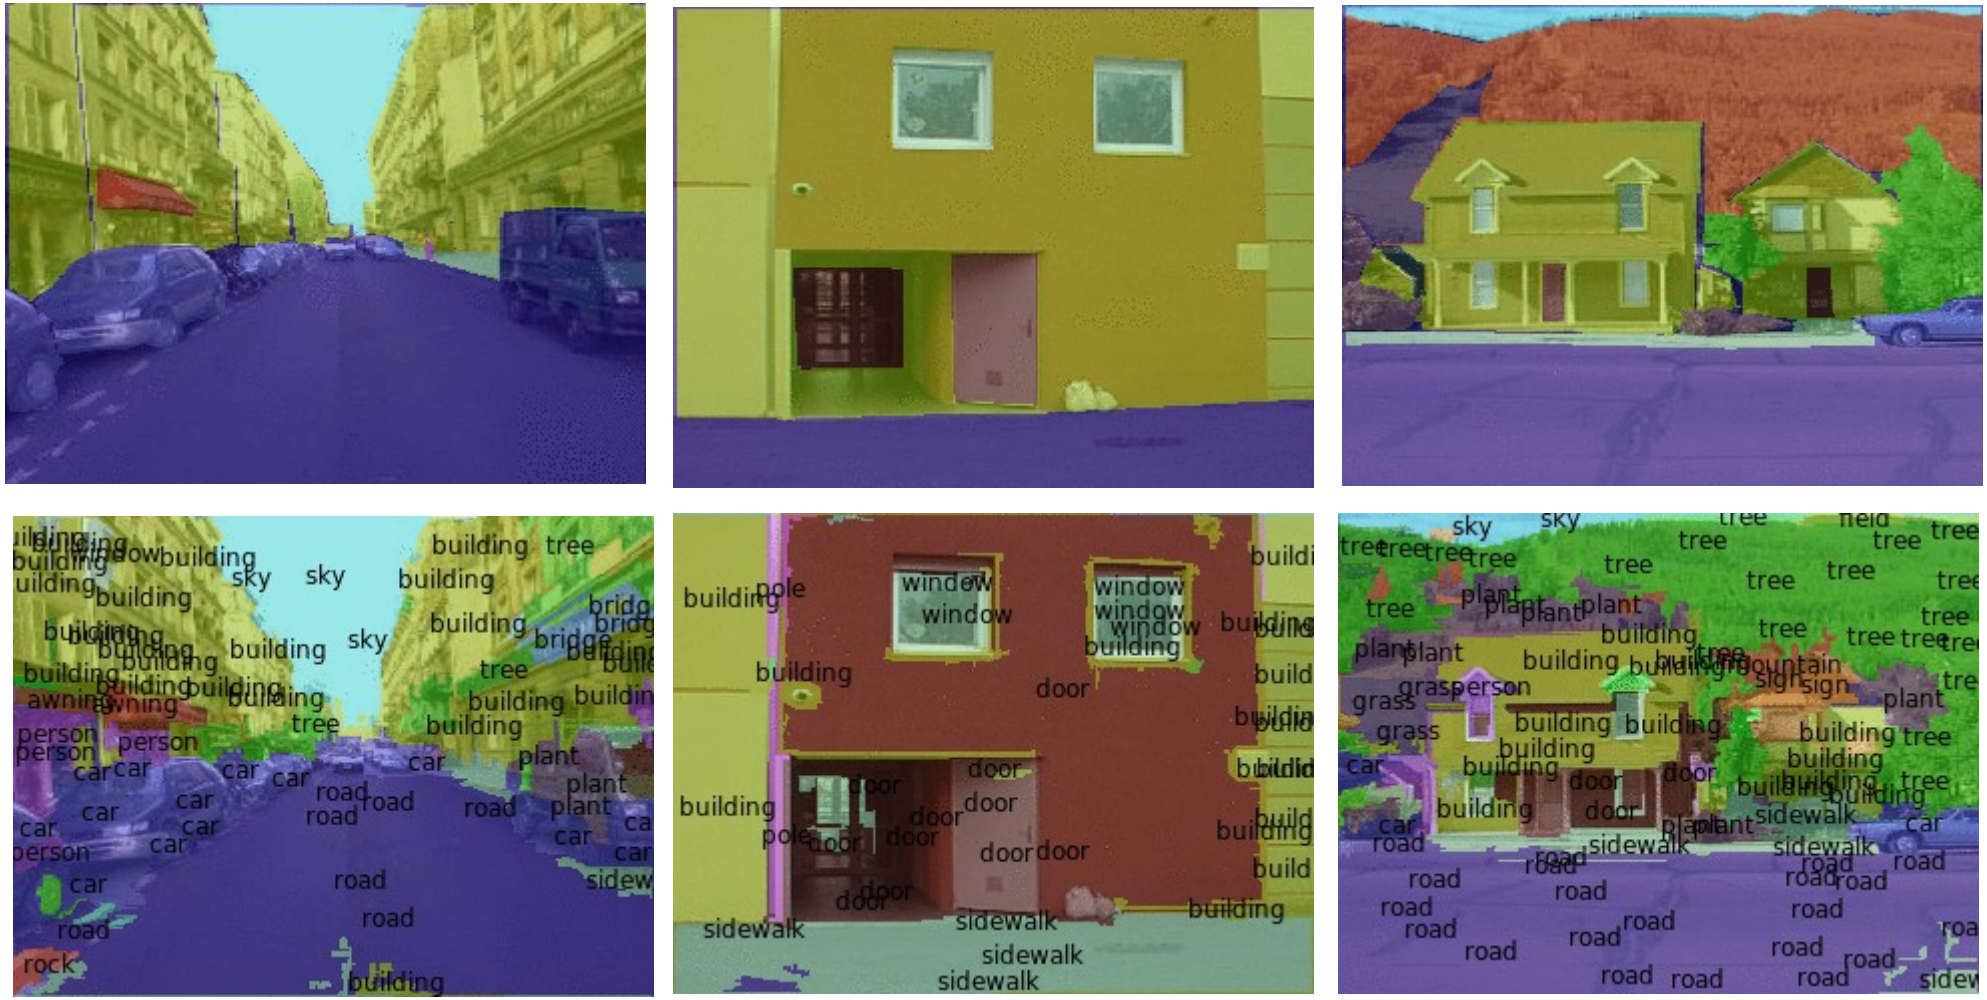

[Farabet et al. ICML 2012, PAMI 2013]

Scene Parsing/Labeling

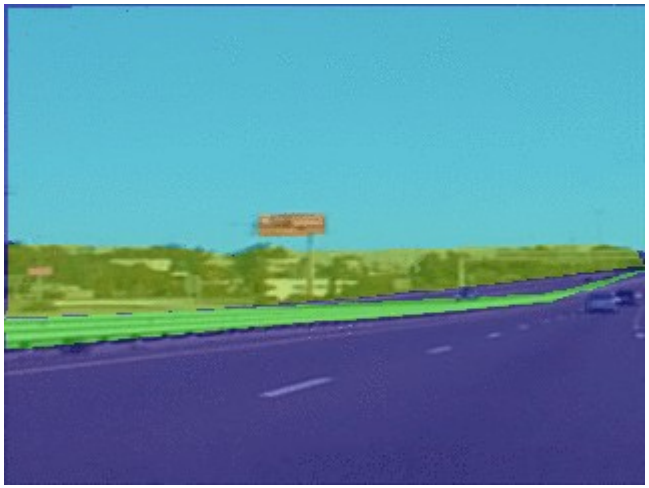

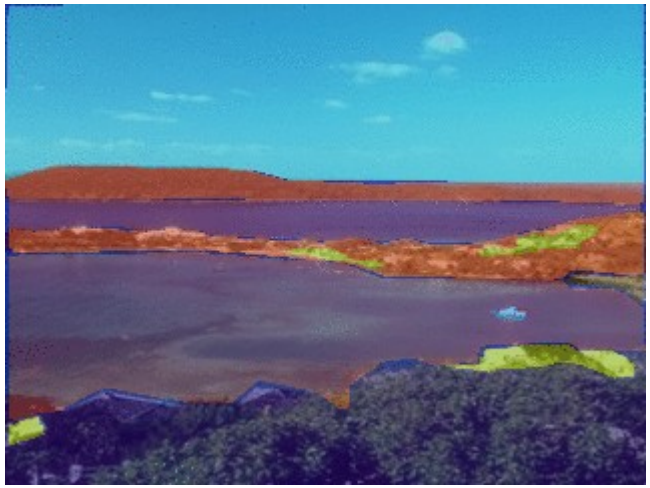

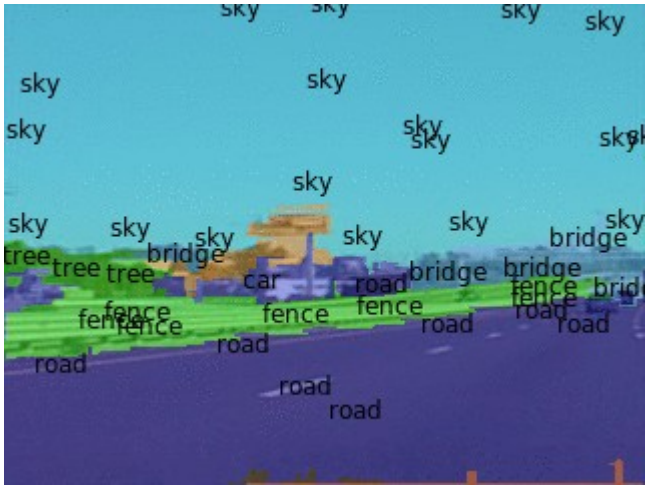

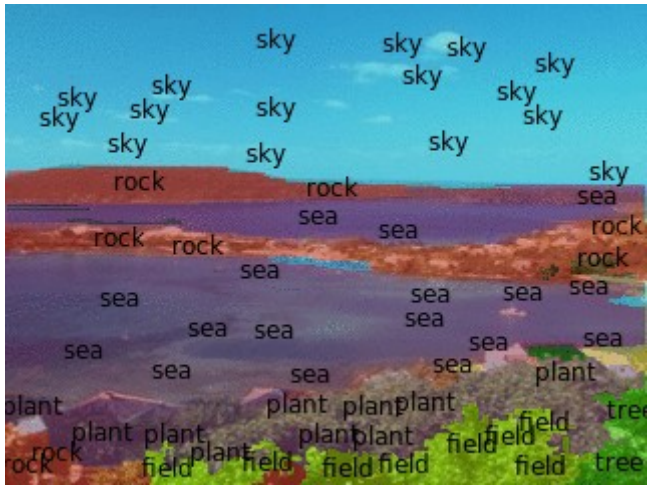

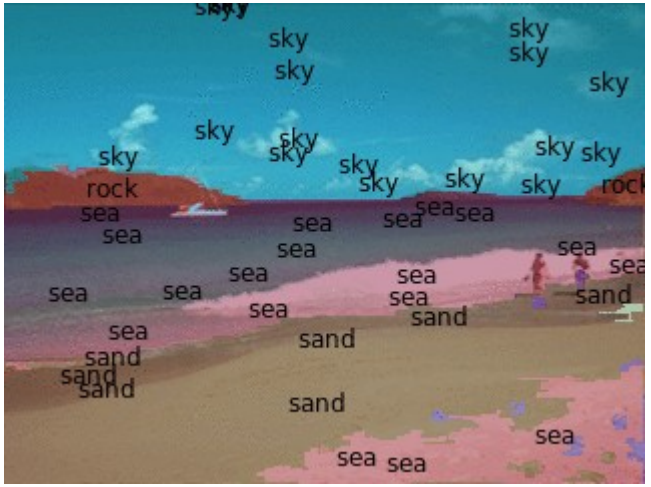

[Farabet et al. ICML 2012, PAMI 2013]

Scene Parsing/Labeling

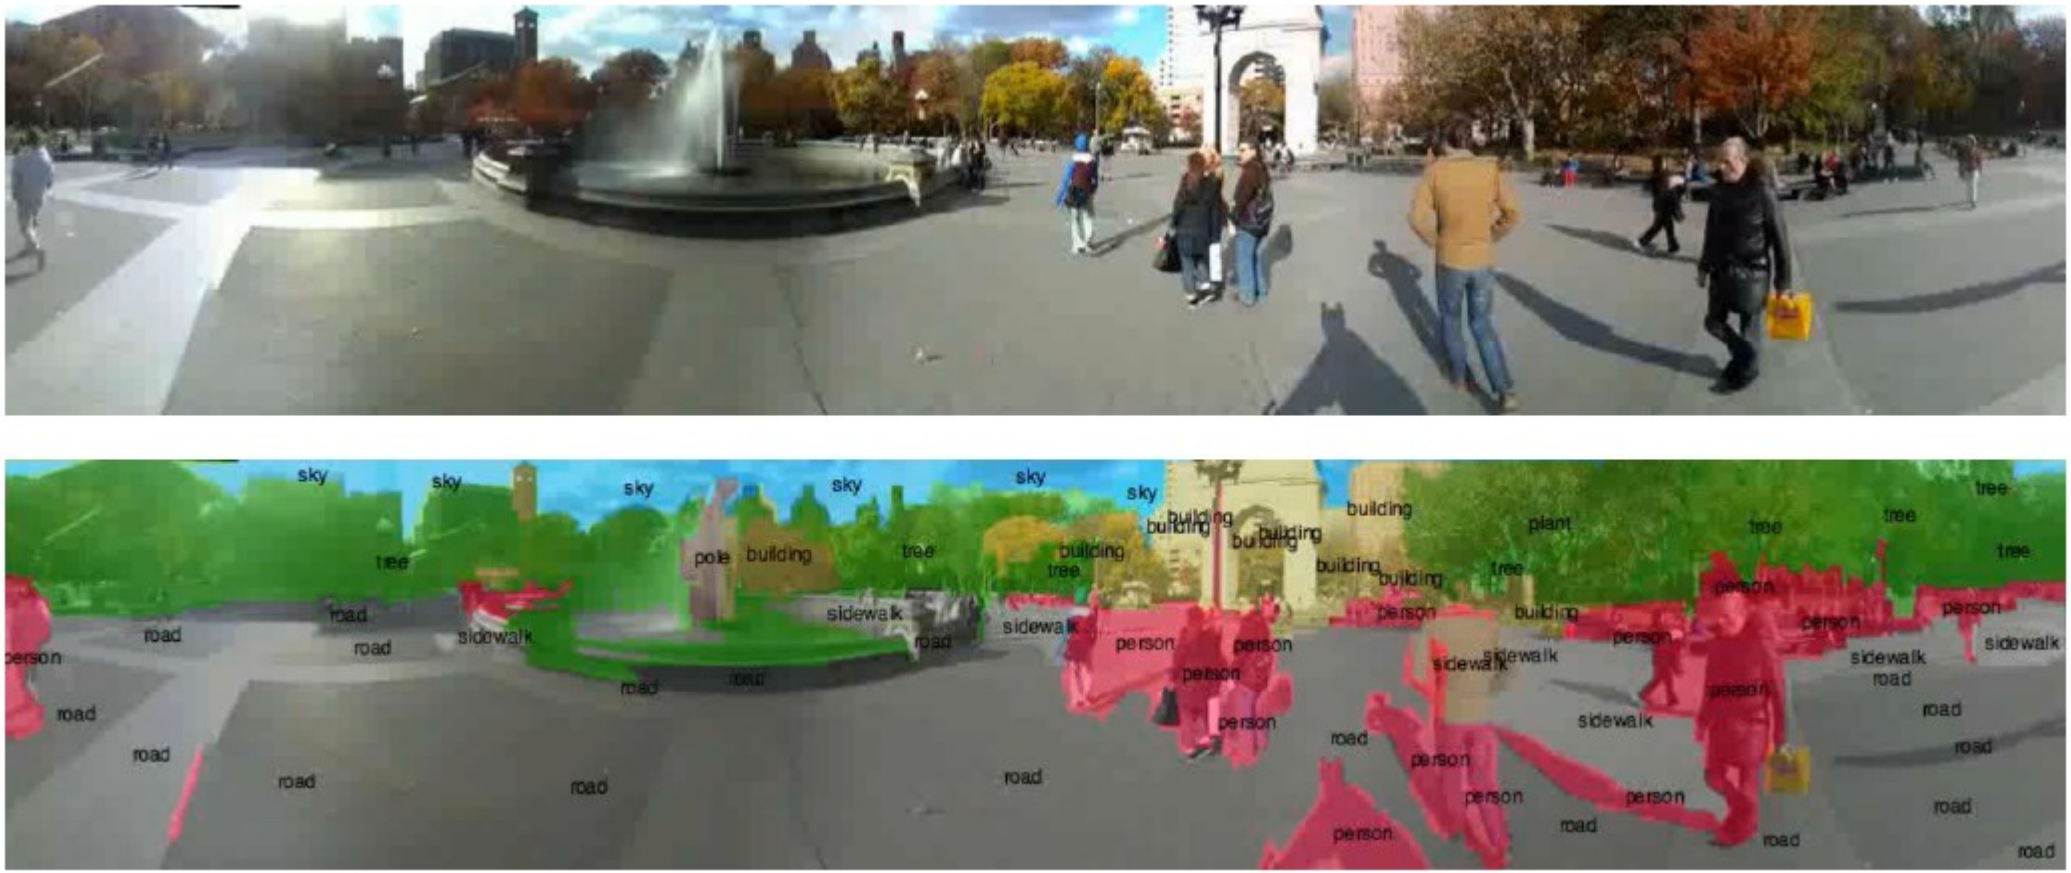

[Farabet et al. ICML 2012, PAMI 2013]

Scene Parsing/Labeling

[Farabet et al. ICML 2012, PAMI 2013]

Scene Parsing/Labeling

- No post-processing
- Frame-by-frame
- ConvNet runs at 50ms/frame on Virtex-6 FPGA hardware
  - But communicating the features over ethernet limits system performance

Temporal Consistency

- Spatio-Temporal Super-Pixel segmentation
  - [Couprie et al ICIP 2013]
  - [Couprie et al JMLR under review]
  - Majority vote over super-pixels

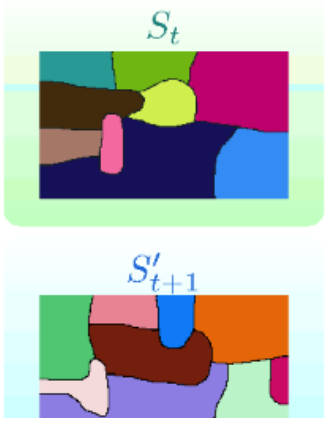

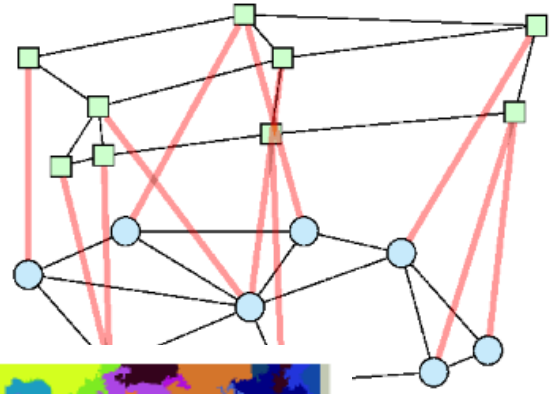

Independent segmentations S'_1, S'_2 and S'_3

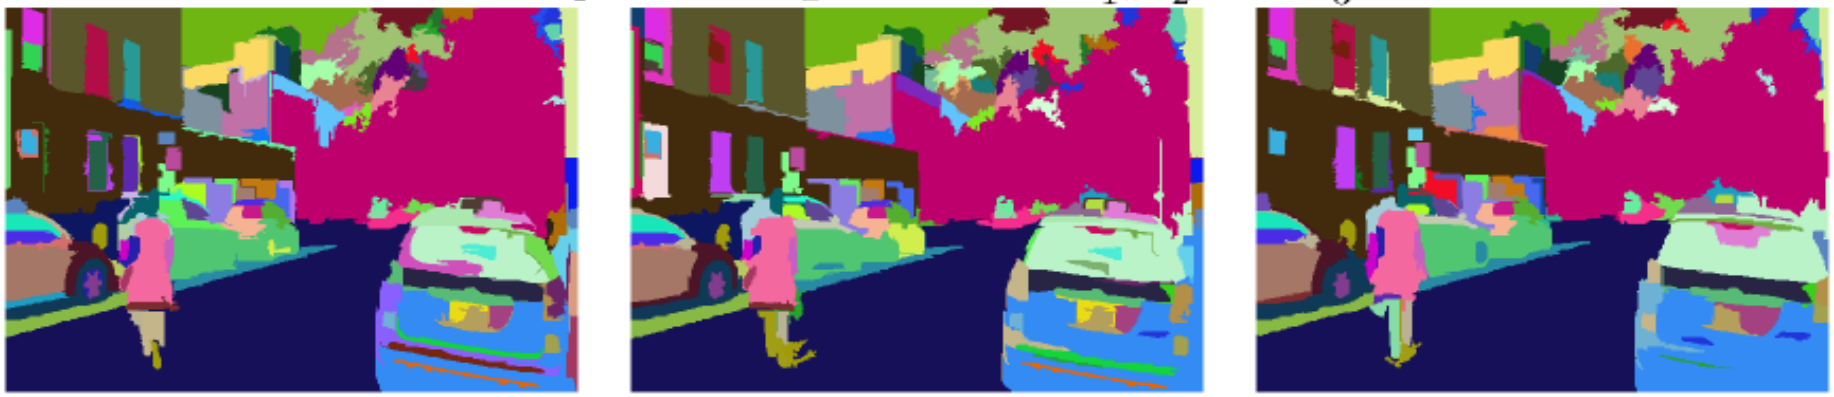

Temporally consistent segmentations $S_1 (= S'_1), S_2$, and S_3

Scene Parsing/Labeling: Temporal Consistency

Y LeCun

- Causal method for temporal consistency

[Couprie, Farabet, Najman, LeCun ICLR 2013, ICIP 2013]

Convolutional Networks
for
"Real" Object Recognition

Then., two things happened...

- **The ImageNet dataset [Fei-Fei et al. 2012]**
  - 1.5 million training samples
  - 1000 categories

- **Fast Graphical Processing Units (GPU)**
  - Capable of 1 trillion operations/second

Matchstick

Flute

Backpack

Sea lion

Strawberry

Bathing cap

Racket

ImageNet Large-Scale Visual Recognition Challenge

The ImageNet dataset

- 1.5 million training samples
- 1000 fine-grained categories (breeds of dogs....)

Object Recognition [Krizhevsky, Sutskever, Hinton 2012]

- Won the 2012 ImageNet LSVRC. 60 Million parameters, 832M MAC ops

4M	FULL CONNECT	4Mflop
16M	FULL 4096/ReLU	16M
37M	FULL 4096/ReLU	37M
	MAX POOLING	
442K	CONV 3x3/ReLU 256fm	74M
1.3M	CONV 3x3ReLU 384fm	224M
884K	CONV 3x3/ReLU 384fm	149M
	MAX POOLING 2x2sub	
	LOCAL CONTRAST NORM	
307K	CONV 11x11/ReLU 256fm	223M
	MAX POOL 2x2sub	
	LOCAL CONTRAST NORM	
35K	CONV 11x11/ReLU 96fm	105M

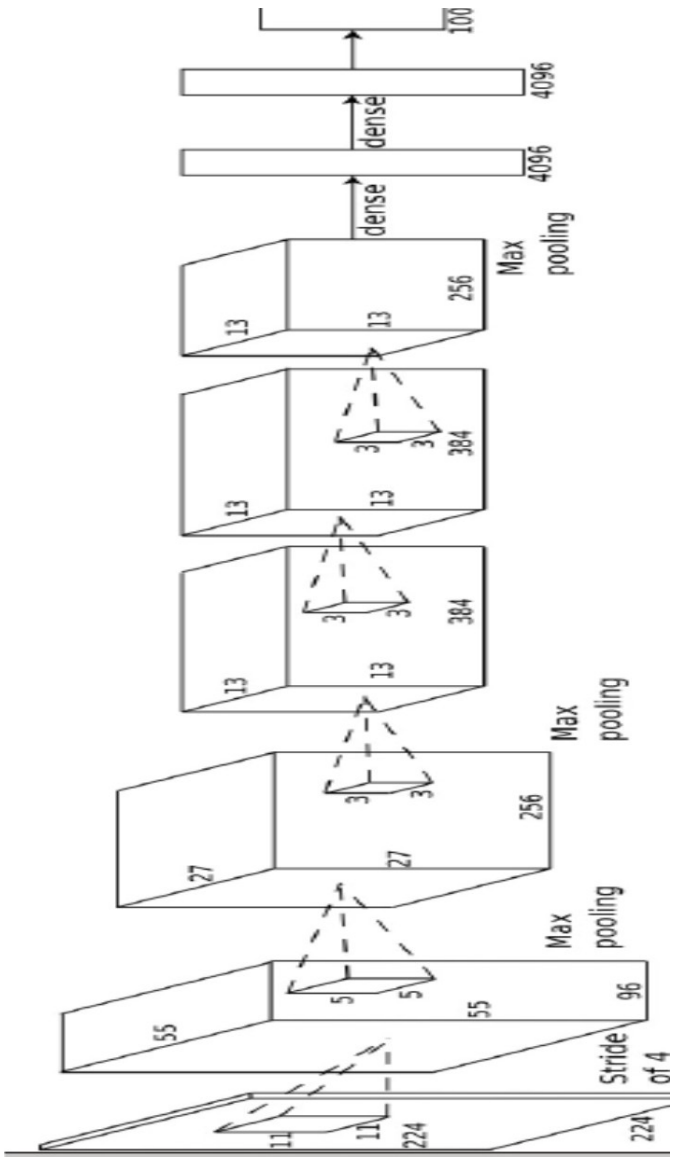

Object Recognition [Krizhevsky, Sutskever, Hinton 2012]

- Method: large convolutional net
  - 650K neurons, 832M synapses, 60M parameters
  - Trained with backprop on GPU
  - Trained “with all the tricks Yann came up with in the last 20 years, plus dropout” (Hinton, NIPS 2012)
  - Rectification, contrast normalization,...
- Error rate: 15% (whenever correct class isn't in top 5)
- Previous state of the art: 25% error

- A REVOLUTION IN COMPUTER VISION

- Acquired by Google in Jan 2013
- Deployed in Google+ Photo Tagging in May 2013

ConvNet-Based Google+ Photo Tagger

Searched my personal collection for "bird"

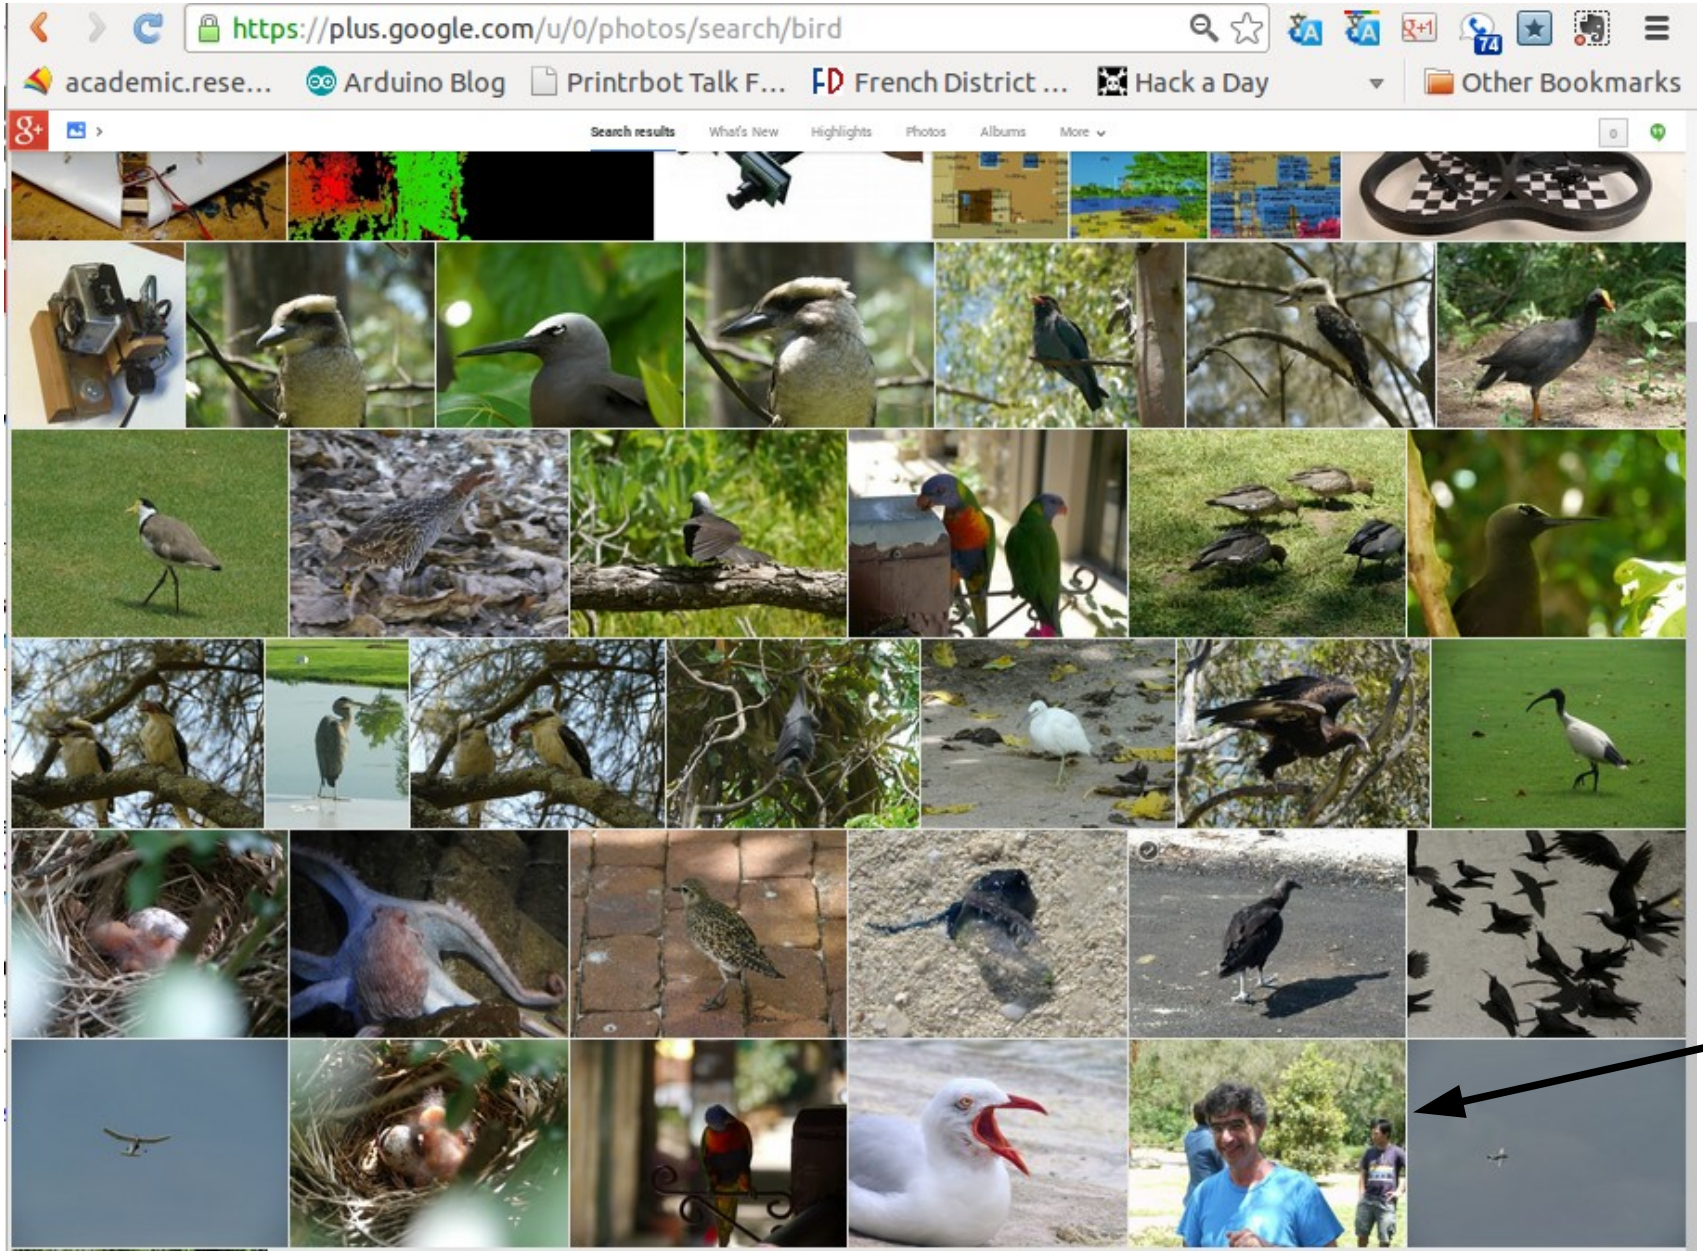

Samy Bengio ???

NYU ConvNet Trained on ImageNet: OverFeat

Y LeCun

- [Sermanet et al. arXiv:1312.6229]
- Trained on GPU using Torch7
- Uses a number of new tricks
- Classification 1000 categories:
  - 13.8% error (top 5) with an ensemble of 7 networks (Krizhevsky: 15%)
  - 15.4% error (top 5) with a single network (Krizhevksy: 18.2%)
- Classification+Localization
  - 30% error (Krizhevsky: 34%)
- Detection (200 categories)
  - 19% correct

- Dowloadable code (running, no training)
  - **Search for "overfeat NYU" on Google**
  - **http://cilvr.nyu.edu → software**

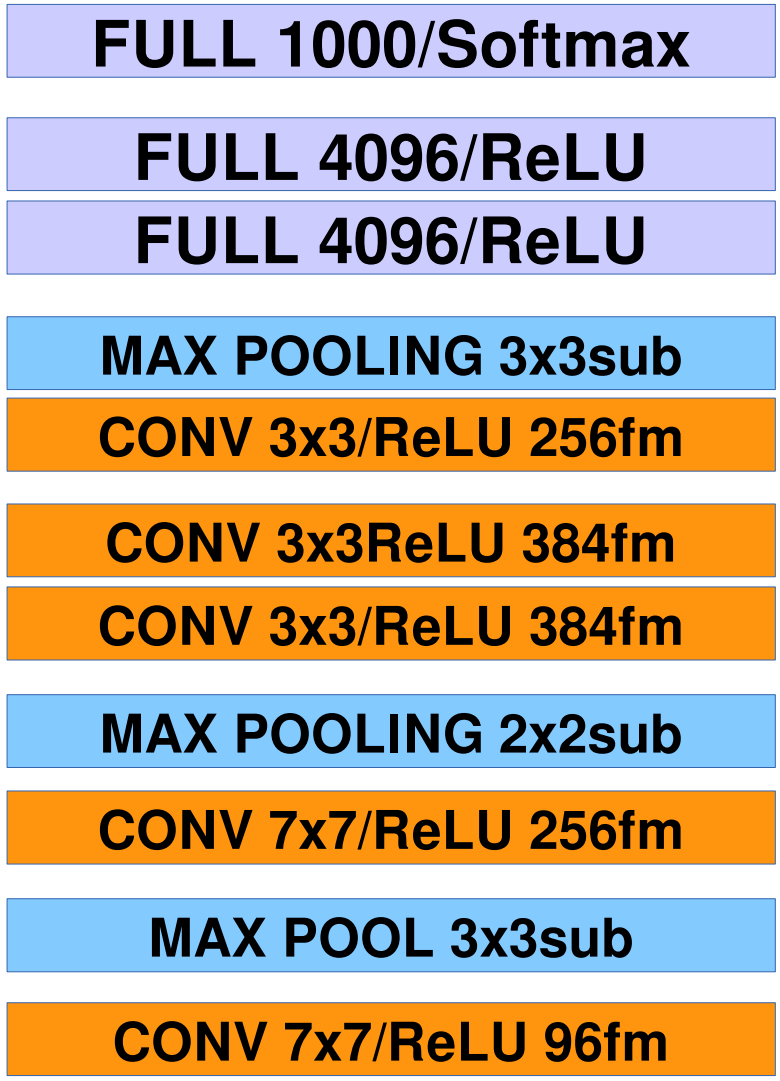

Kernels: Layer 1 (7x7)and Layer 2 (7x7)

- Layer 1: 3x96 kernels, RGB->96 feature maps, 7x7 Kernels, stride 2

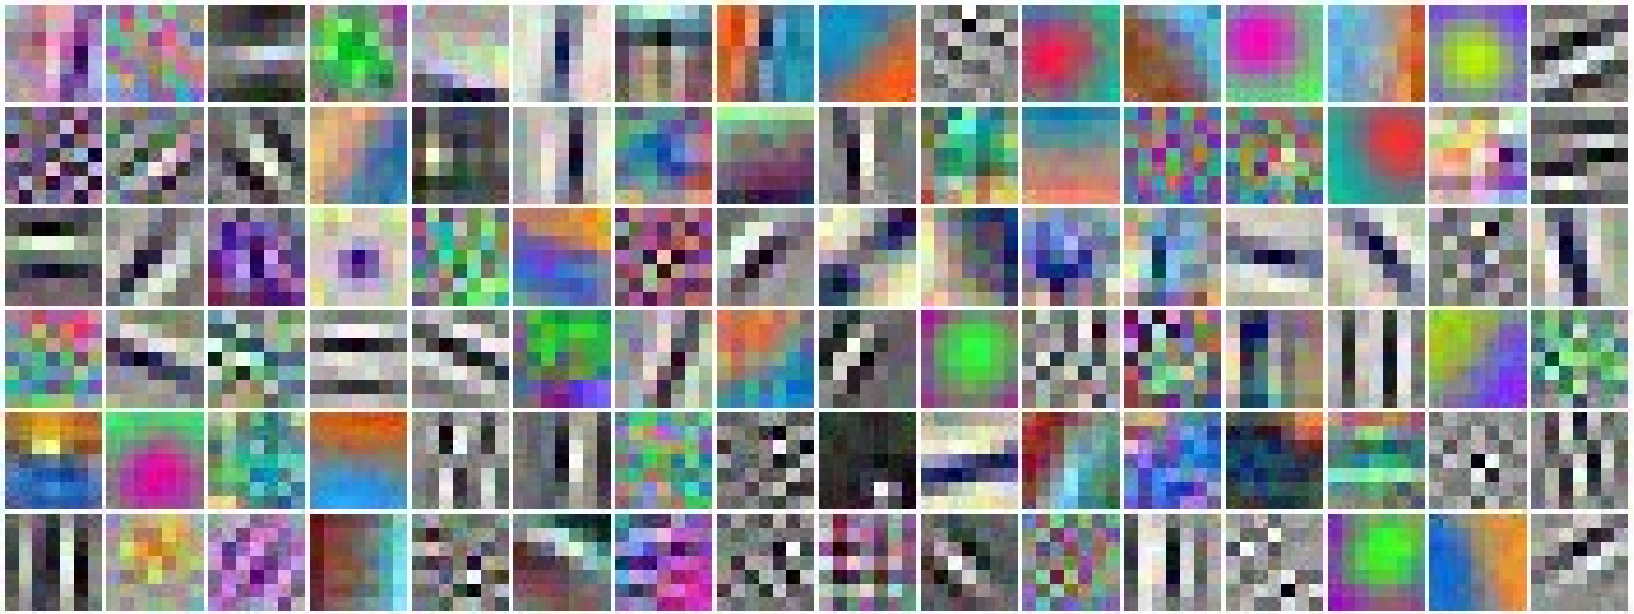

- Layer 2: 96x256 kernels, 7x7

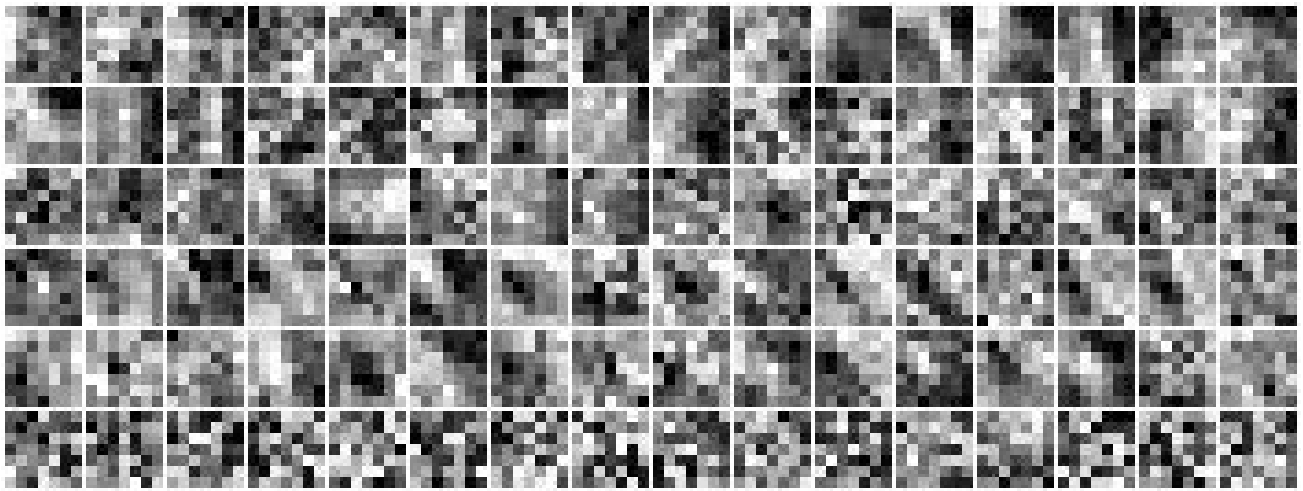

Kernels: Layer 1 (11x11)

- Layer 1: 3x96 kernels, RGB->96 feature maps, 11x11 Kernels, stride 4

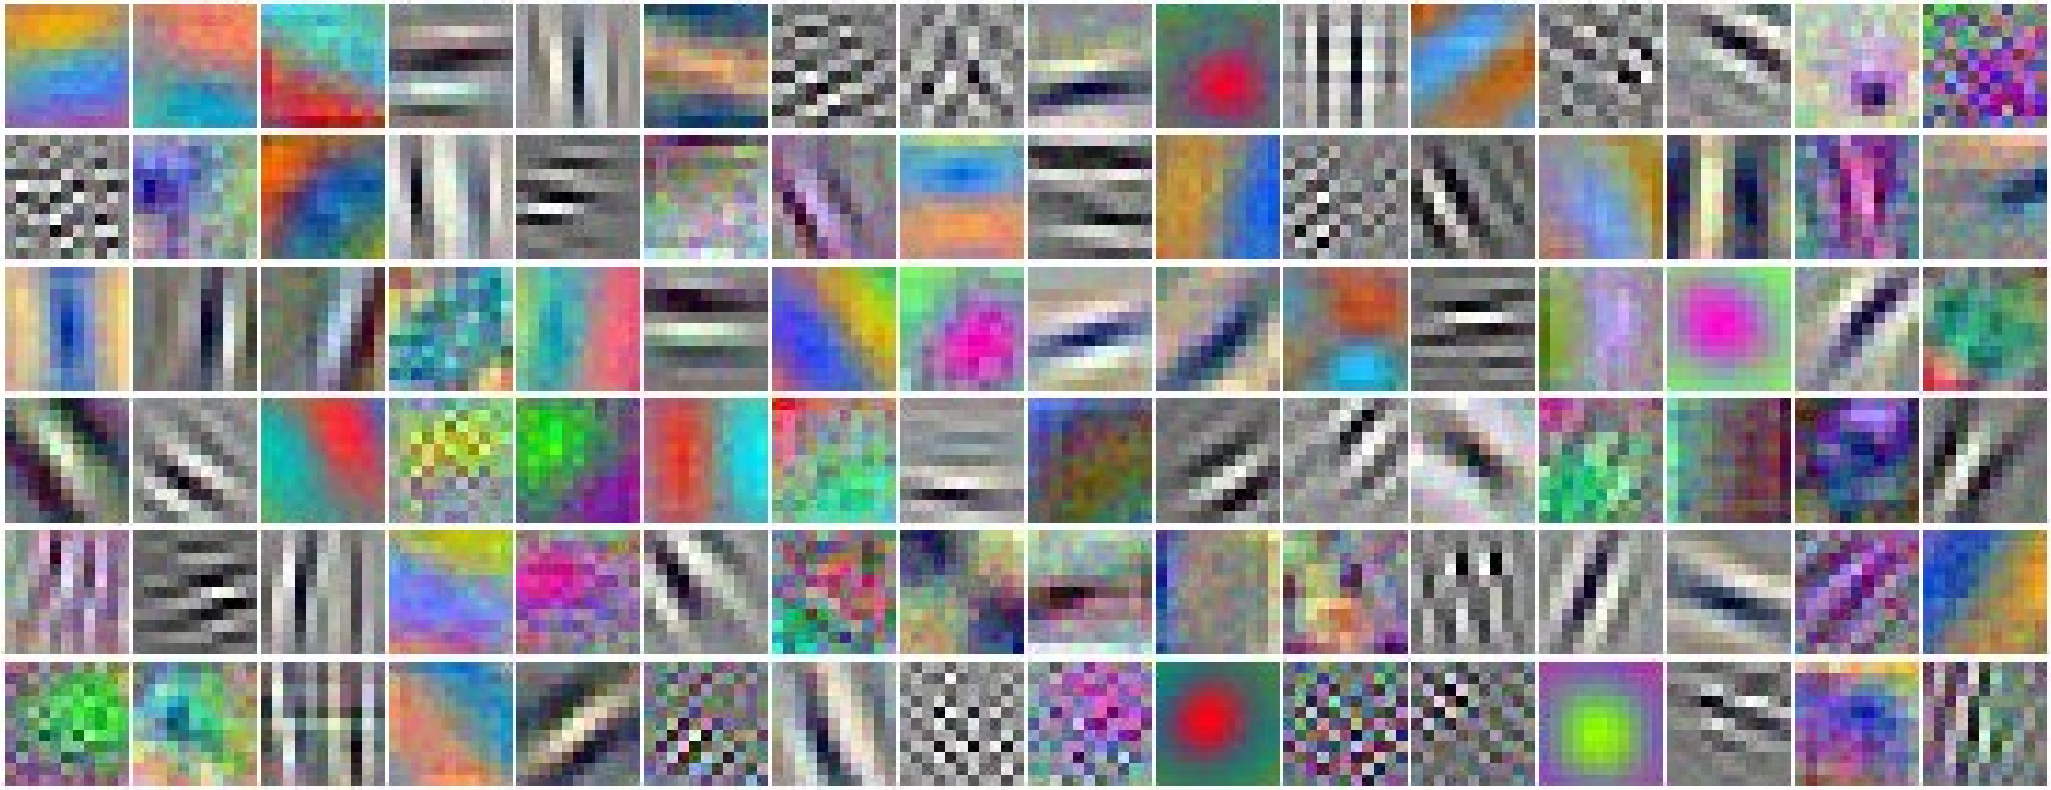

ImageNet: Classification

- Give the name of the dominant object in the image
- Top-5 error rates: if correct class is not in top 5, count as error
  - Red:ConvNet, blue: no ConvNet

2012 Teams	%error
Supervision (Toronto)	15.3
ISI (Tokyo)	26.1
VGG (Oxford)	26.9
XRCE/INRIA	27.0
UvA (Amsterdam)	29.6
INRIA/LEAR	33.4

2013 Teams	%error
Clarifai (NYU spinoff)	11.7
NUS (singapore)	12.9
Zeiler-Fergus (NYU)	13.5
A. Howard	13.5
OverFeat (NYU)	14.1
UvA (Amsterdam)	14.2
Adobe	15.2
VGG (Oxford)	15.2
VGG (Oxford)	23.0

2014 Teams	%error
GoogLeNet	6.6
VGG (Oxford)	7.3
MSRA	8.0
A. Howard	8.1
DeeperVision	9.5
NUS-BST	9.7
TTIC-ECP	10.2
XYZ	11.2
UvA	12.1

Classification+Localization. Results

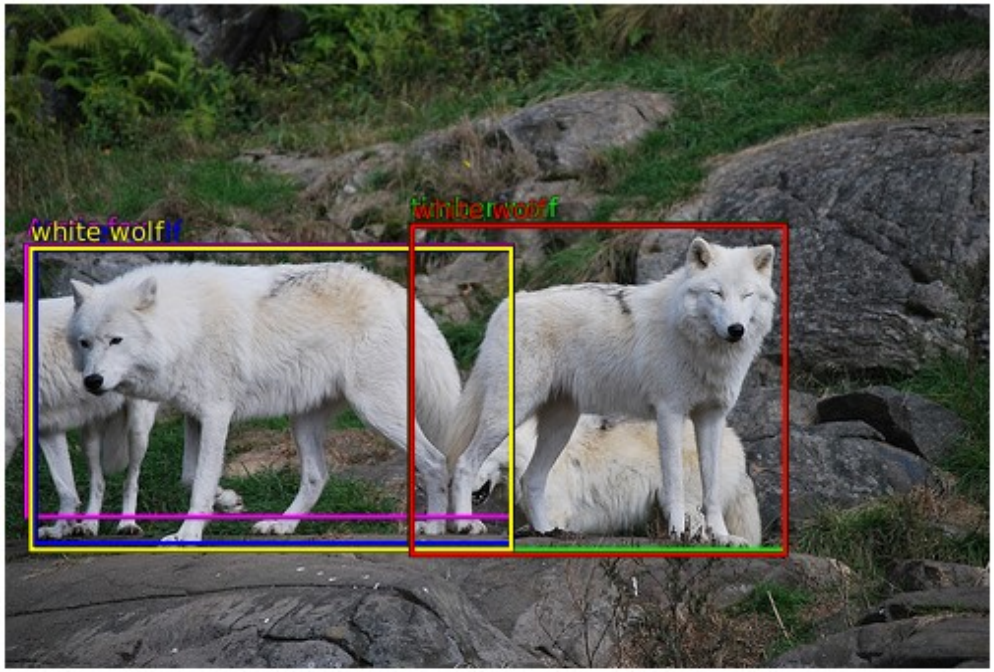

Top 5:
white wolf
white wolf
timber wolf
timber wolf
Arctic fox

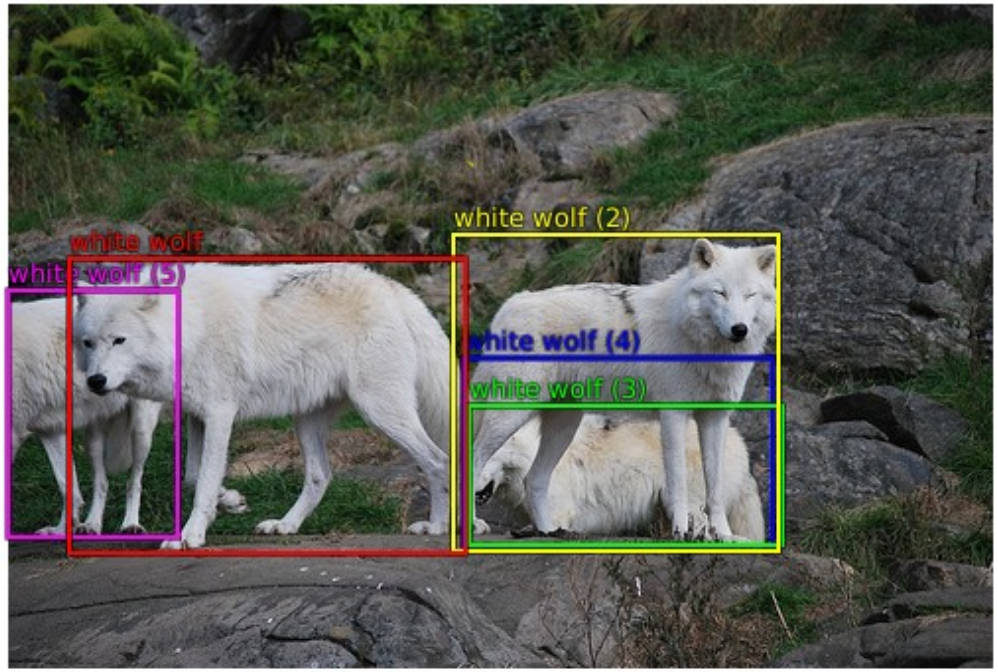

Groundtruth:
white wolf
white wolf (2)
white wolf (3)
white wolf (4)
white wolf (5)

ILSVRC2012_val_00000027.JPEG

Classification+Localization. Error Rates

- It's best to propose several categories for the same window
  - One of them might be right
- 2014 results: 25.3% (VGG Oxford), 26.4% (GoogLeNet)

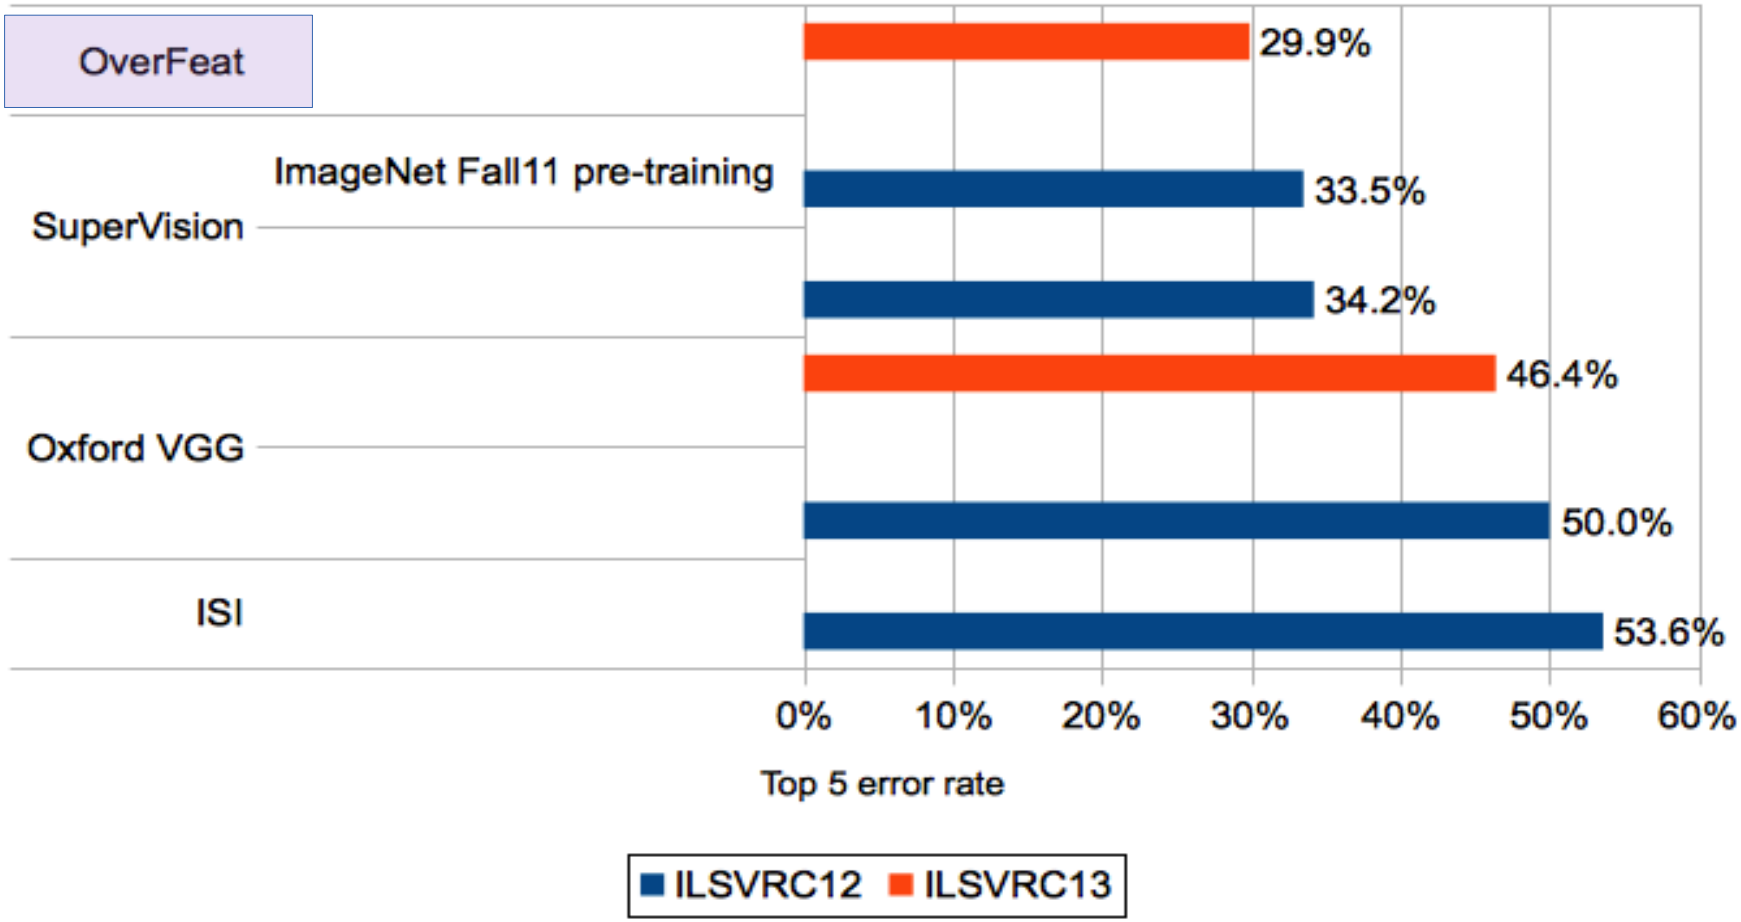

Classification + Localization: multiscale sliding window

Y LeCun

- Apply convnet with a sliding window over the image at multiple scales
- Important note: it's very cheap to slide a convnet over an image
  - Just compute the convolutions over the whole image and replicate the fully-connected layers

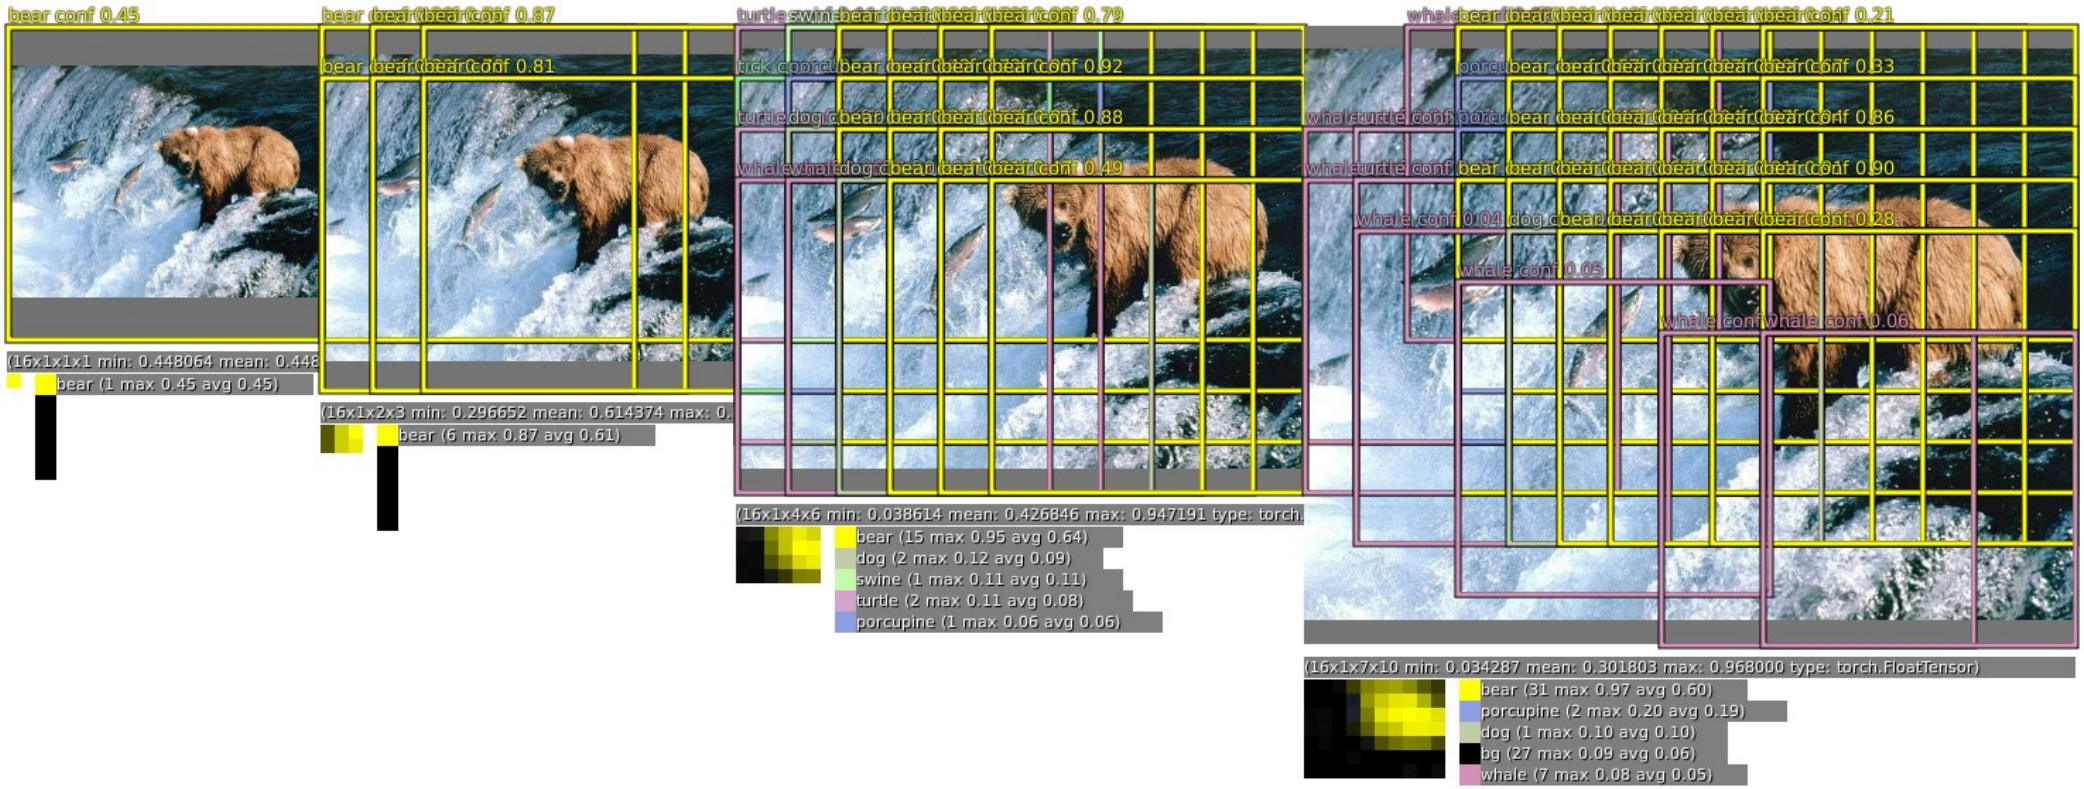

Applying a ConvNet on Sliding Windows is Very Cheap!

Y LeCun

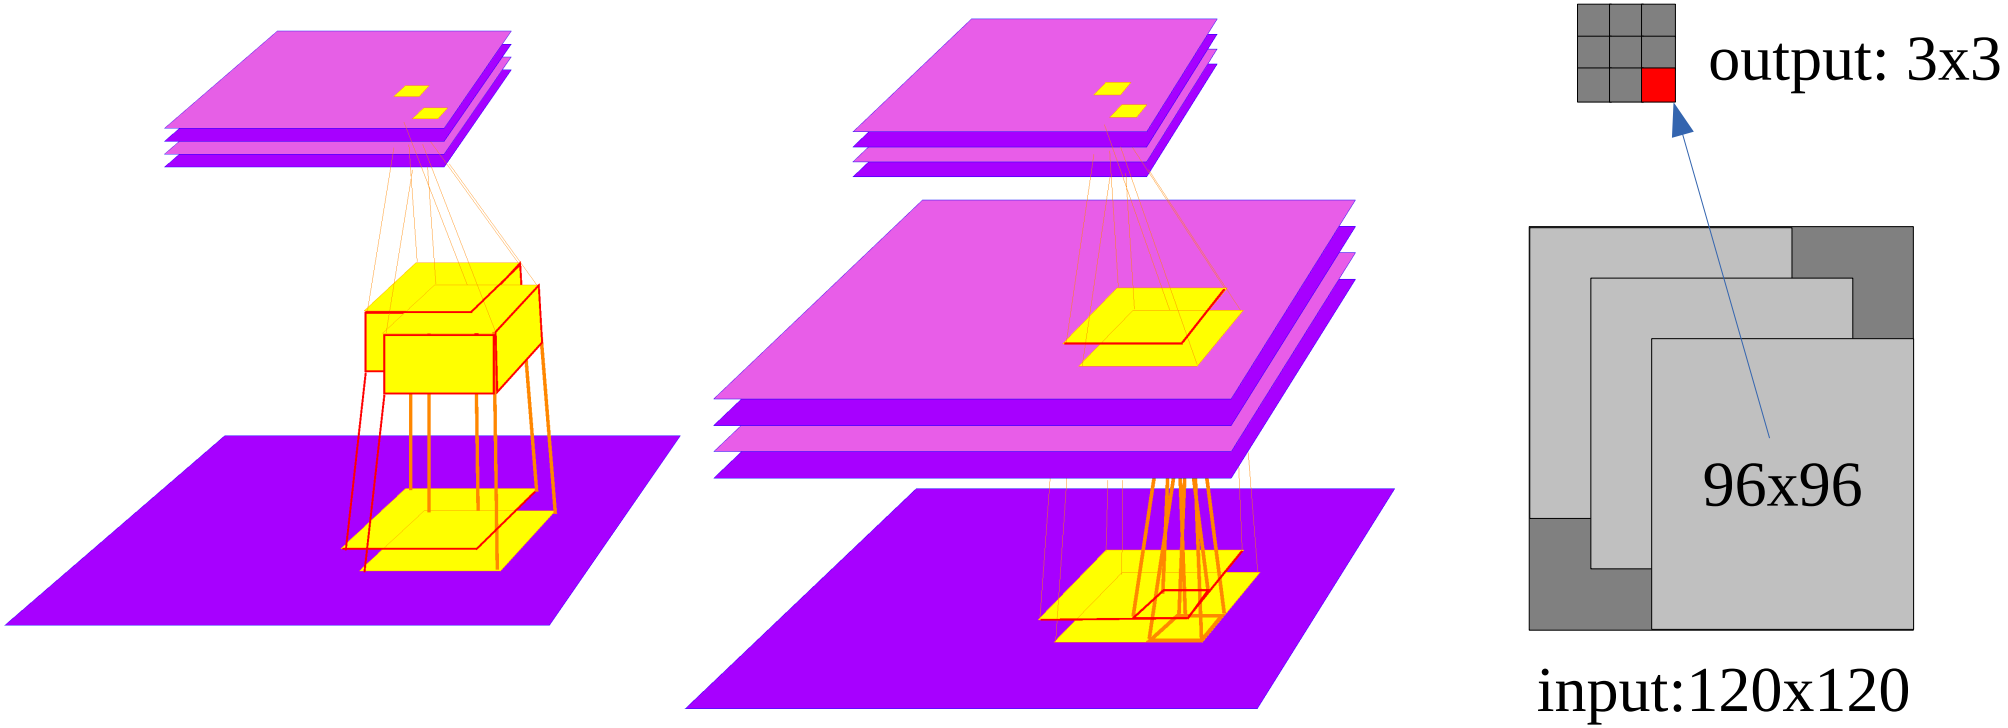

- Traditional Detectors/Classifiers must be applied to every location on a large input image, at multiple scales.
- Convolutional nets can replicated over large images very cheaply.
- Simply apply the convolutions to the entire image and spatially replicate the fully-connected layers

Classification + Localization: sliding window + bounding box regression

Y LeCun

- Apply convnet with a sliding window over the image at multiple scales
- For each window, predict a class and bounding box parameters
  - Evenif the object is not completely contained in the viewing window, the convnet can predict where it thinks the object is.

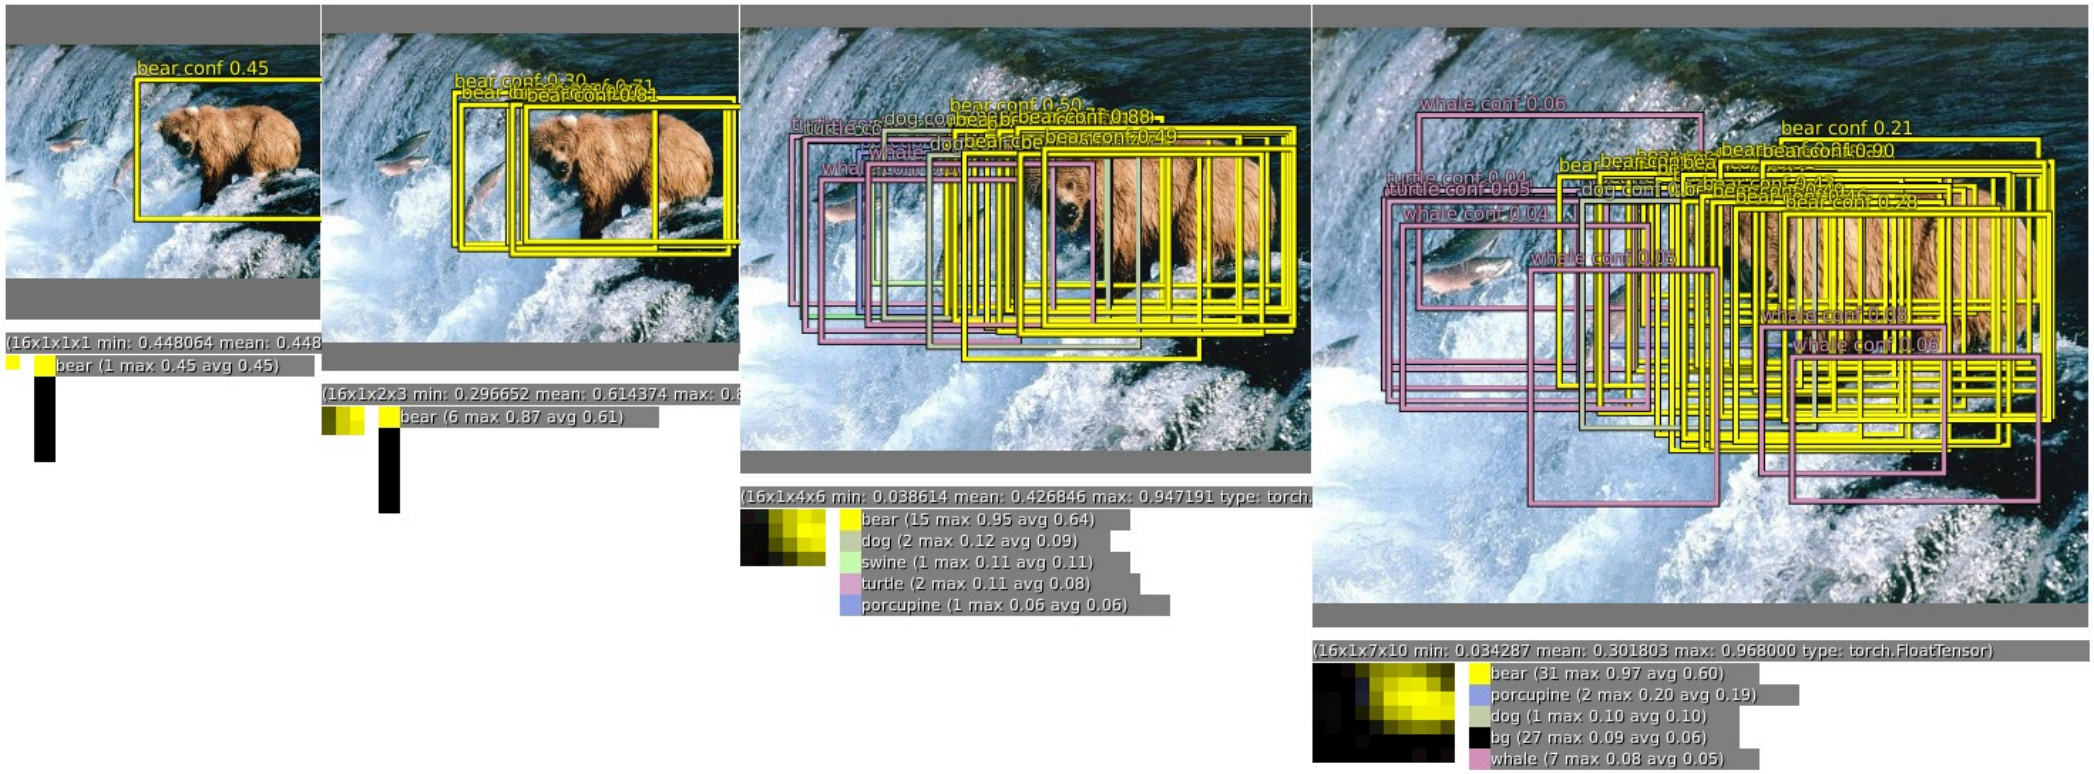

Classification + Localization: sliding window + bounding box regression + bbox voting

Y LeCun

- Apply convnet with a sliding window over the image at multiple scales
- For each window, predict a class and bounding box parameters
- Compute an “average” bounding box, weighted by scores

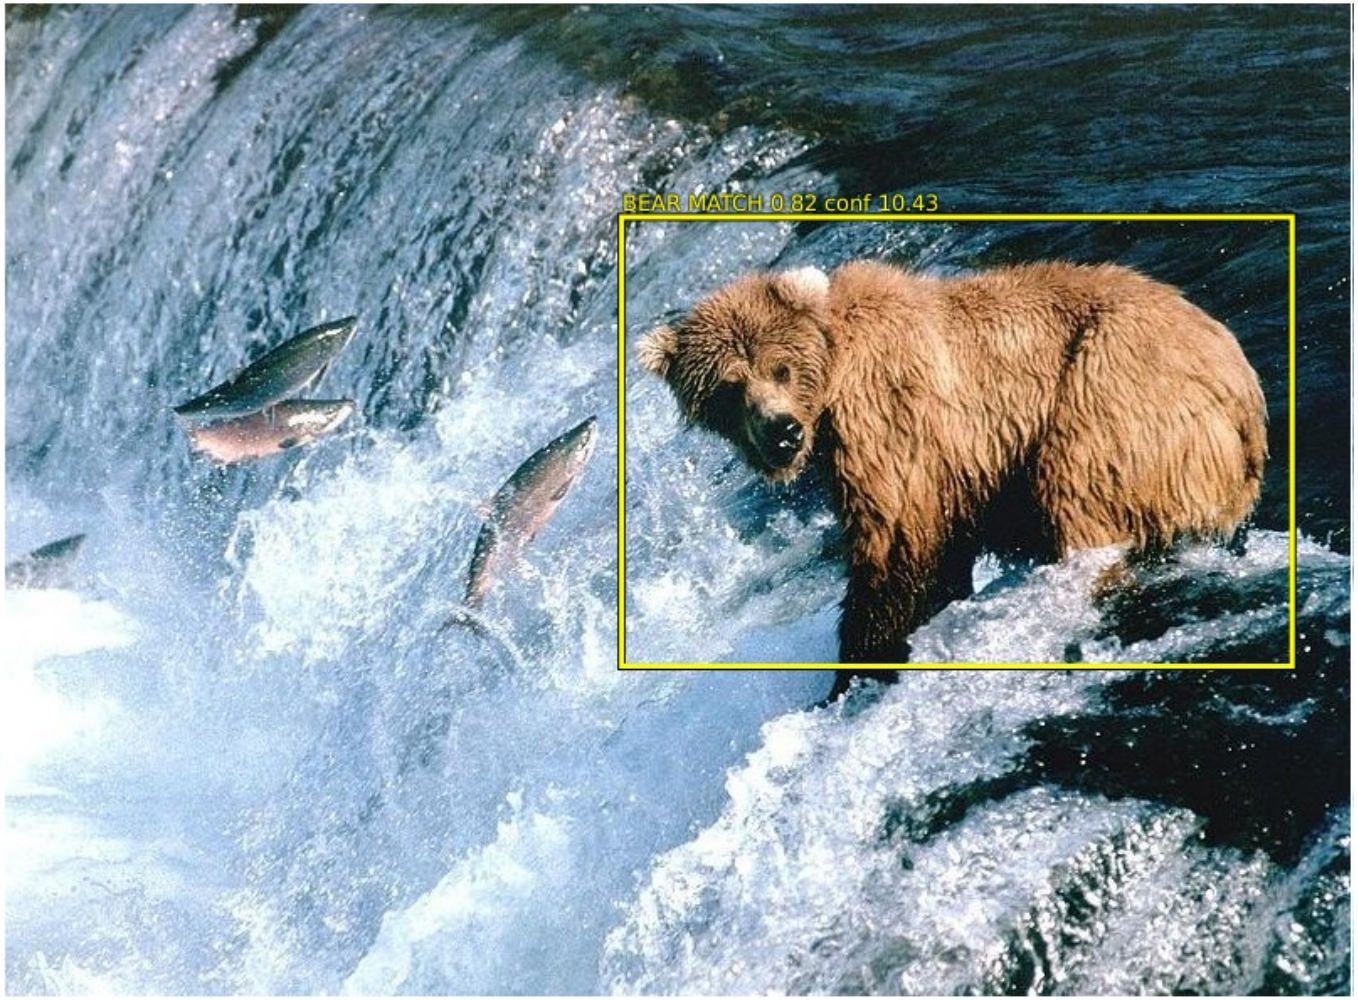

Localization: Sliding Window + bbox vote + multiscale

Y LeCun

Detection / Localization

Detection / Localization

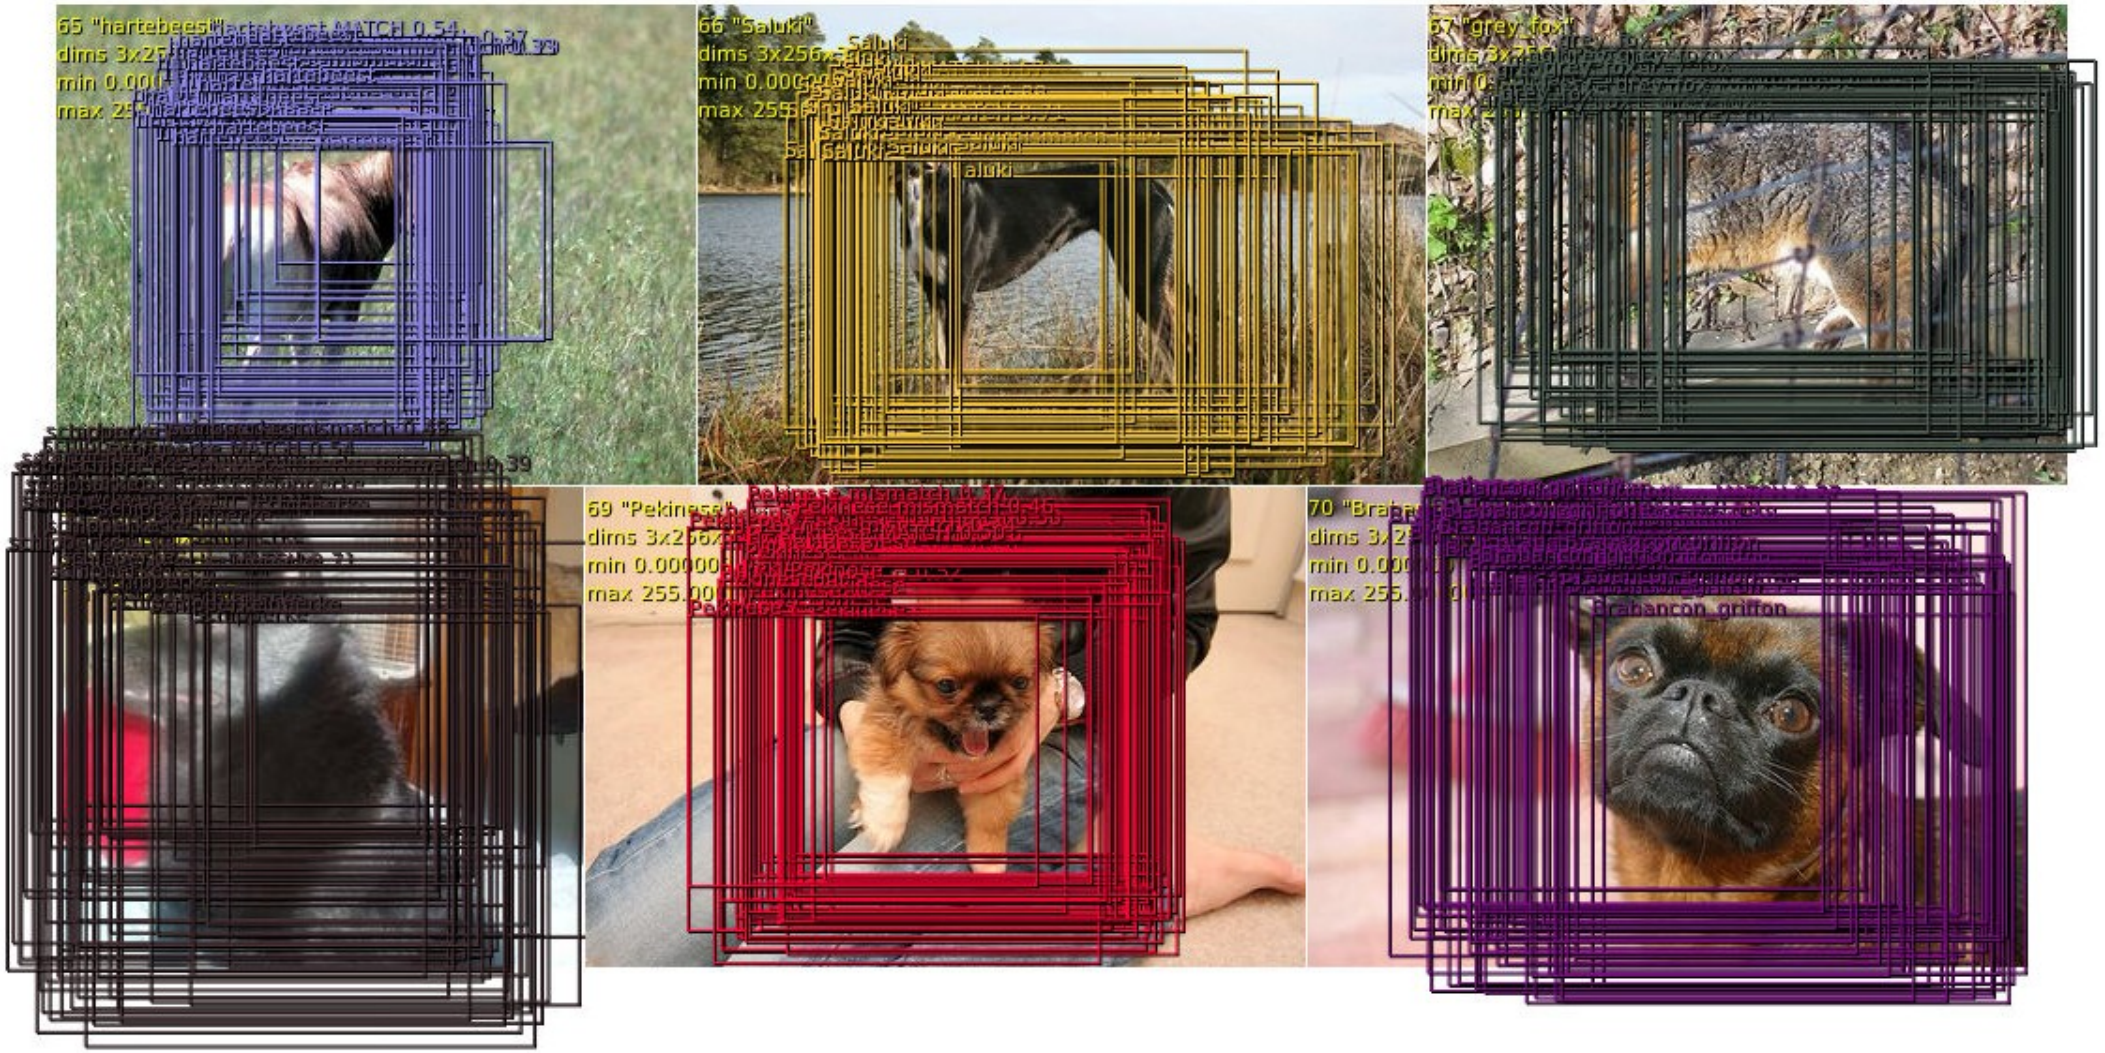

Detection / Localization

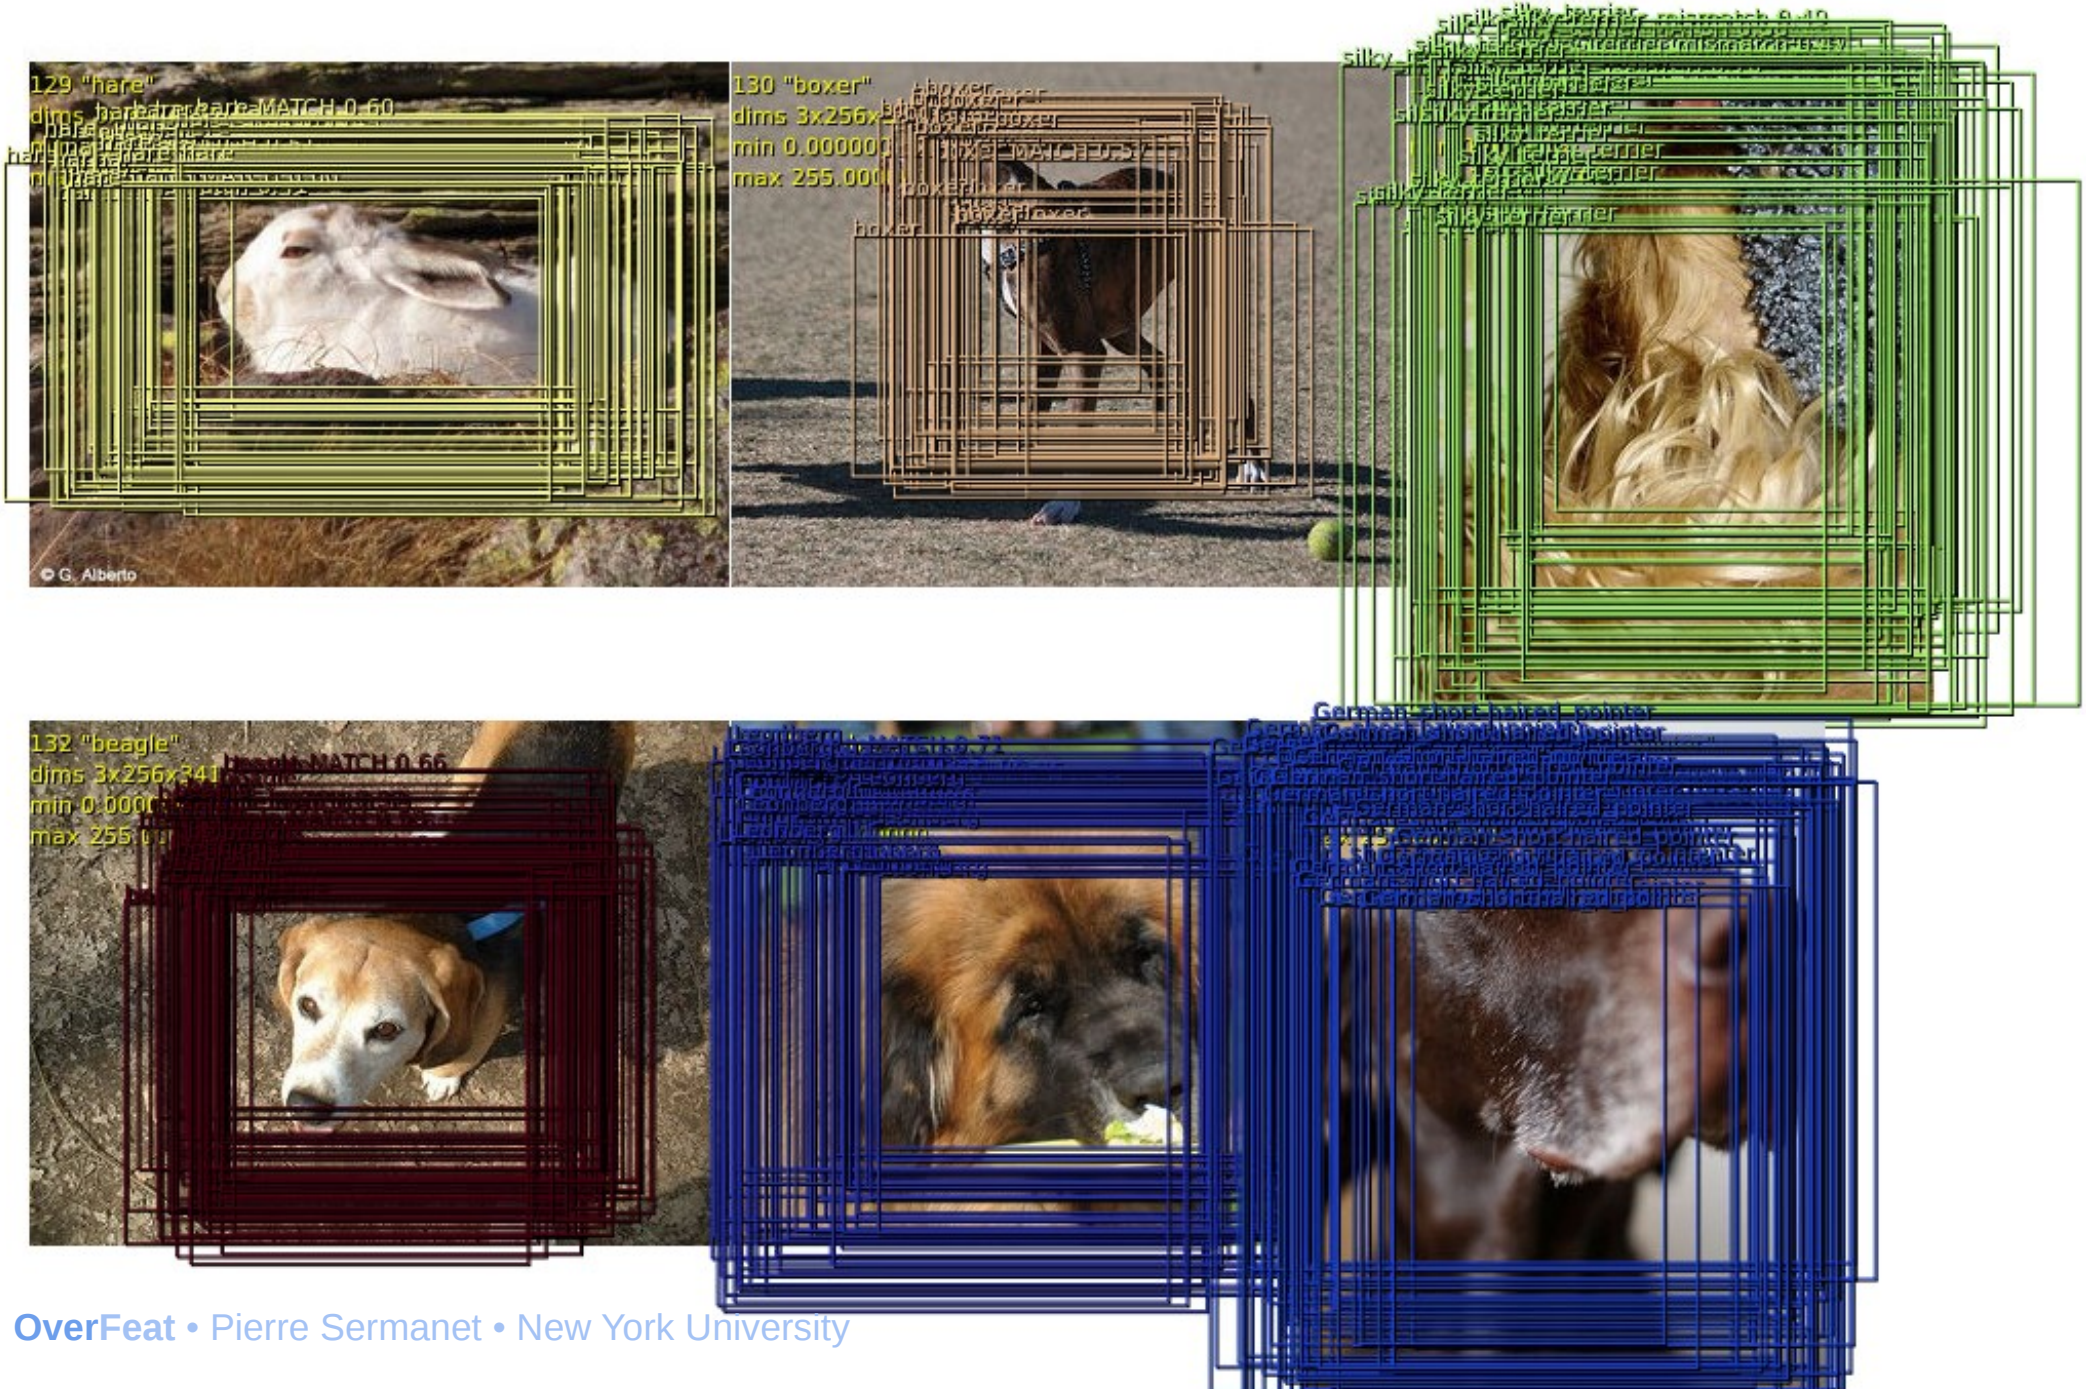

Detection / Localization

Y LeCun

Detection / Localization

Y LeCun

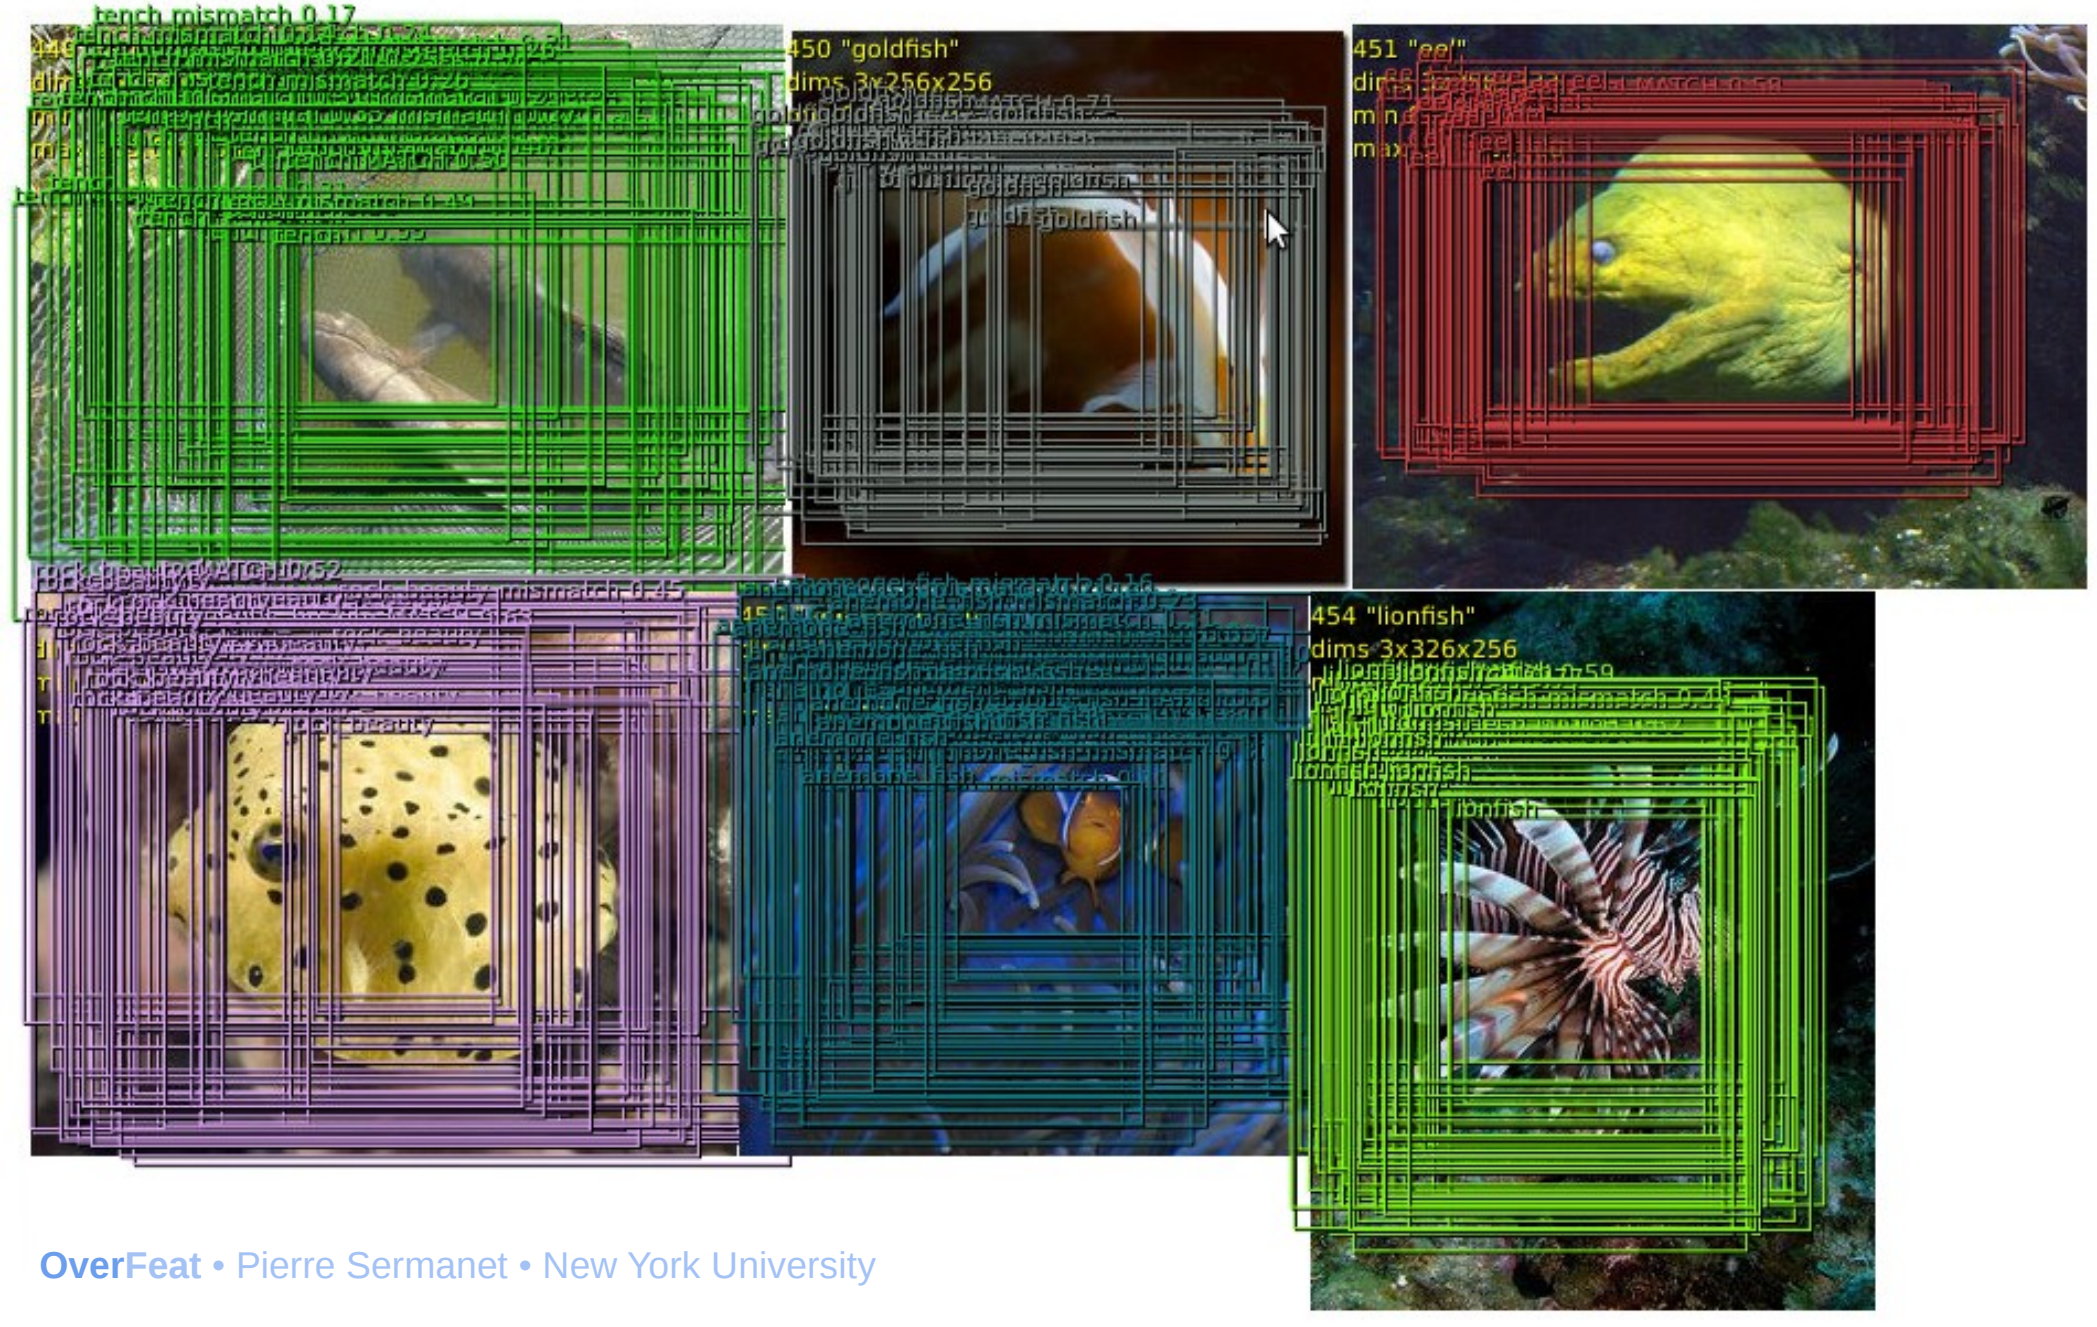

Detection: Examples

- 200 broad categories
- There is a penalty for false positives
- Some examples are easy some are impossible/ambiguous
- Some classes are well detected
  - Burritos?

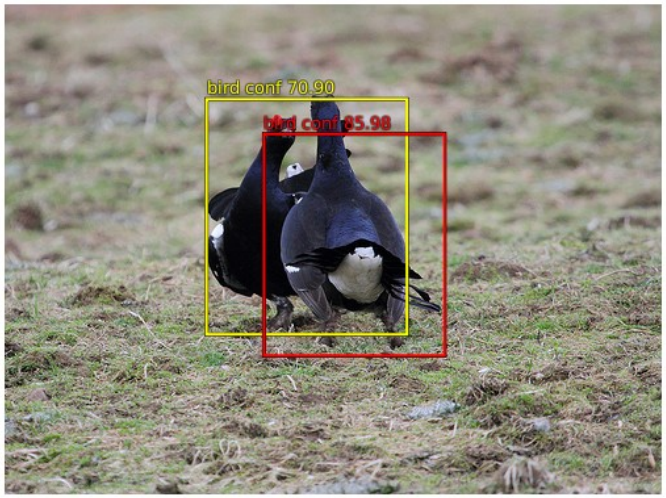

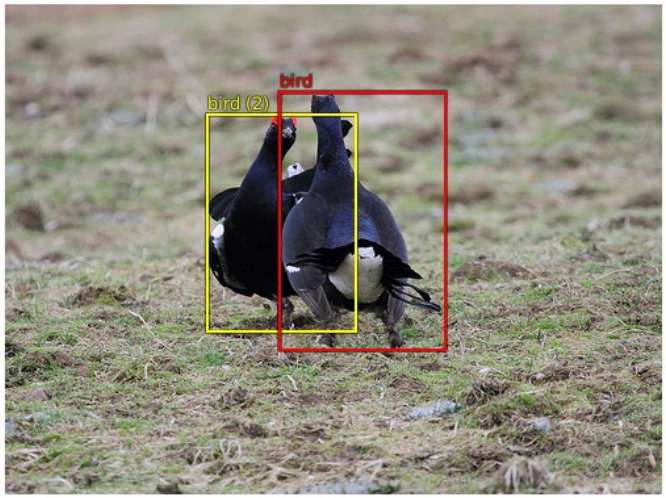

Top predictions:
bird (confidence 86.(
bird (confidence 70.9
ILSVRC2012_val_00001136.JPEG

Groundtruth:

Top predictions:
burrito (confidence 28.9)
ILSVRC2012_val_00000572.JPEG

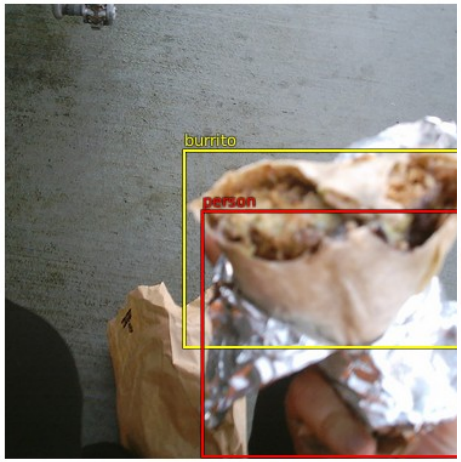

Groundtruth:
person
burrito

Top predictions:
burrito (confidence 17.4)
ILSVRC2012_val_00000606.JPEG

Groundtruth:
burrito
burrito (2)

Detection: Examples

- Groundtruth is sometimes ambiguous or incomplete
- Large overlap between objects stops non-max suppression from working

Top predictions:
tv or monitor (confidence 11.5)
person (confidence 4.5)
miniskirt (confidence 3.1)

ILSVRC2012_val_00000119.JPEG

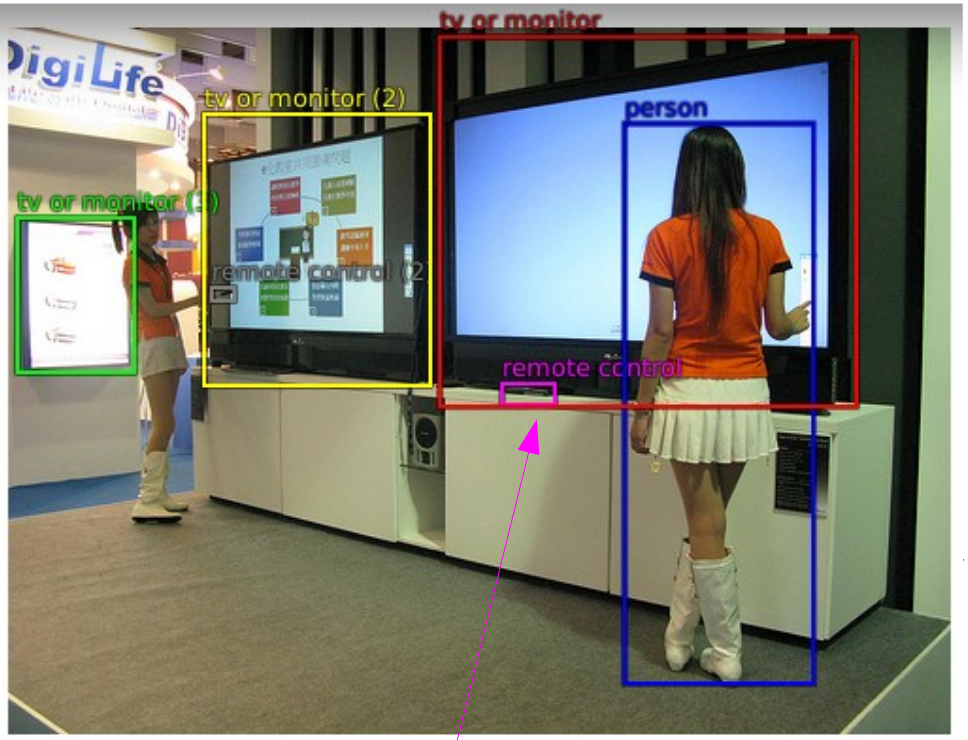

Groundtruth:
tv or monitor
tv or monitor (2)
tv or monitor (3)
person
remote control
remote control (2)

ImageNet: Detection (200 categories)

- Give a bounding box and a category for all objects in the image
- MAP = mean average precision
  - Red:ConvNet, blue: no ConvNet

2013 Teams	mAP
UvA Amsterdam	22.6
NEC Labs-America	20.9
OverFeat NYU	19.4

Off cycle results	mAP
Berkeley RCNN	34.5
OverFeat NYU	24.3

2014 Teams	mAP
GoogLeNet	43.9
CUHK-DeepID2	40.7
DeepInsight	40.4
NUS	37.2
UvA Amsterdam	35.4
MSRA	35.1
Berkeley RCNN	34.5

Results: pre-trained on ImageNet1K, fine-tuned on ImageNet Detection

Y LeCun

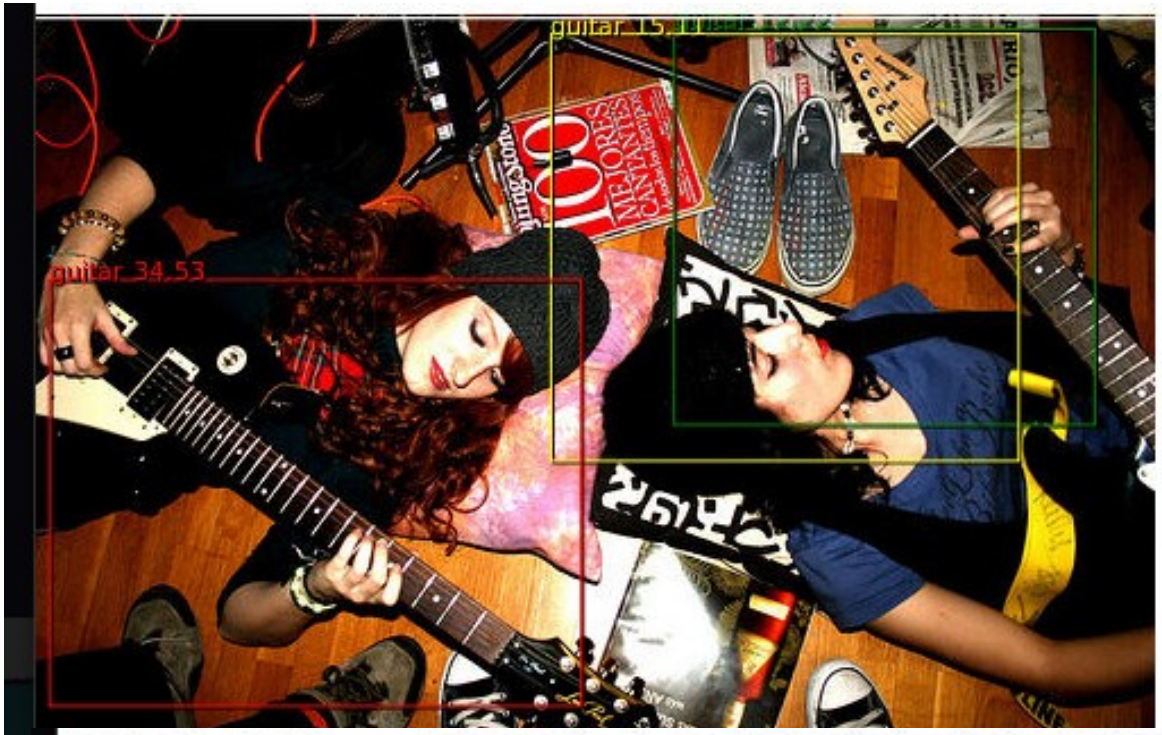

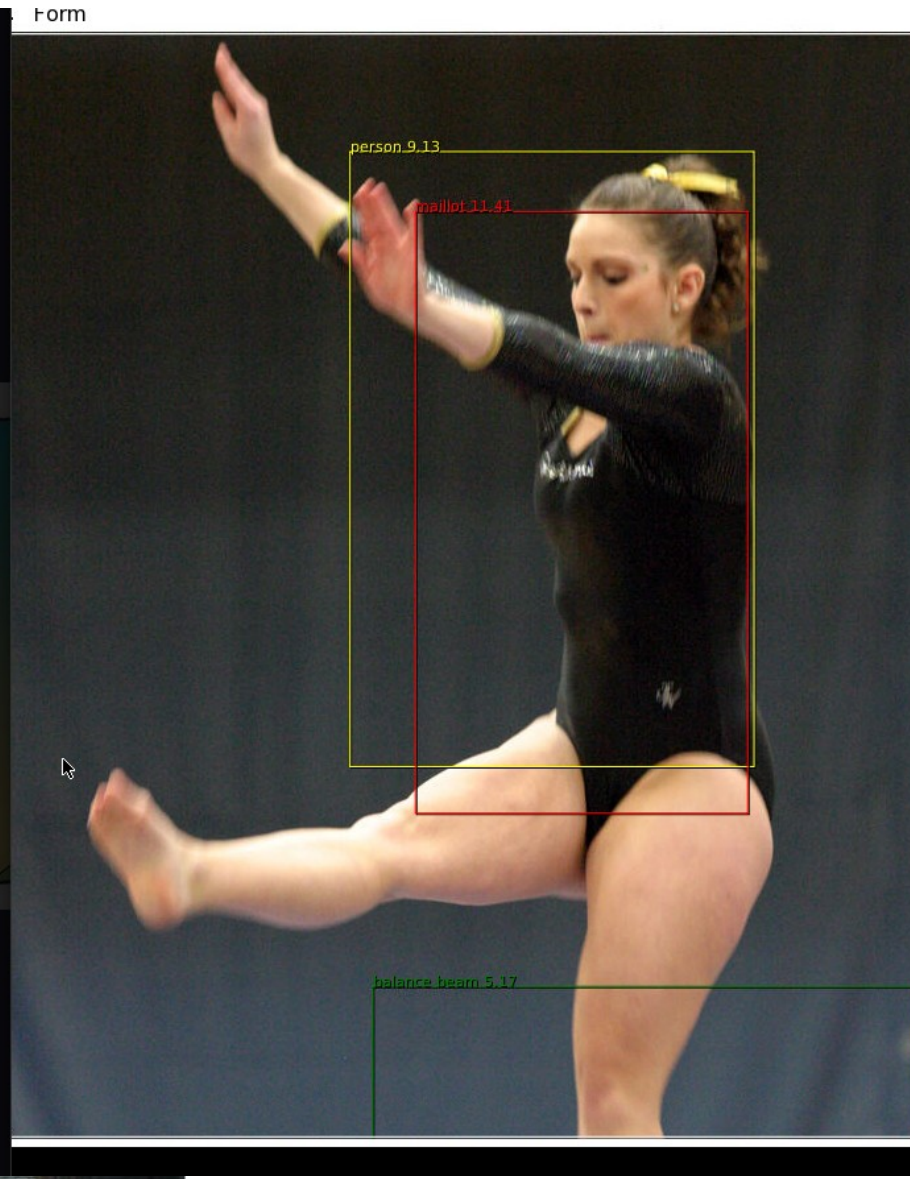

Detection Examples

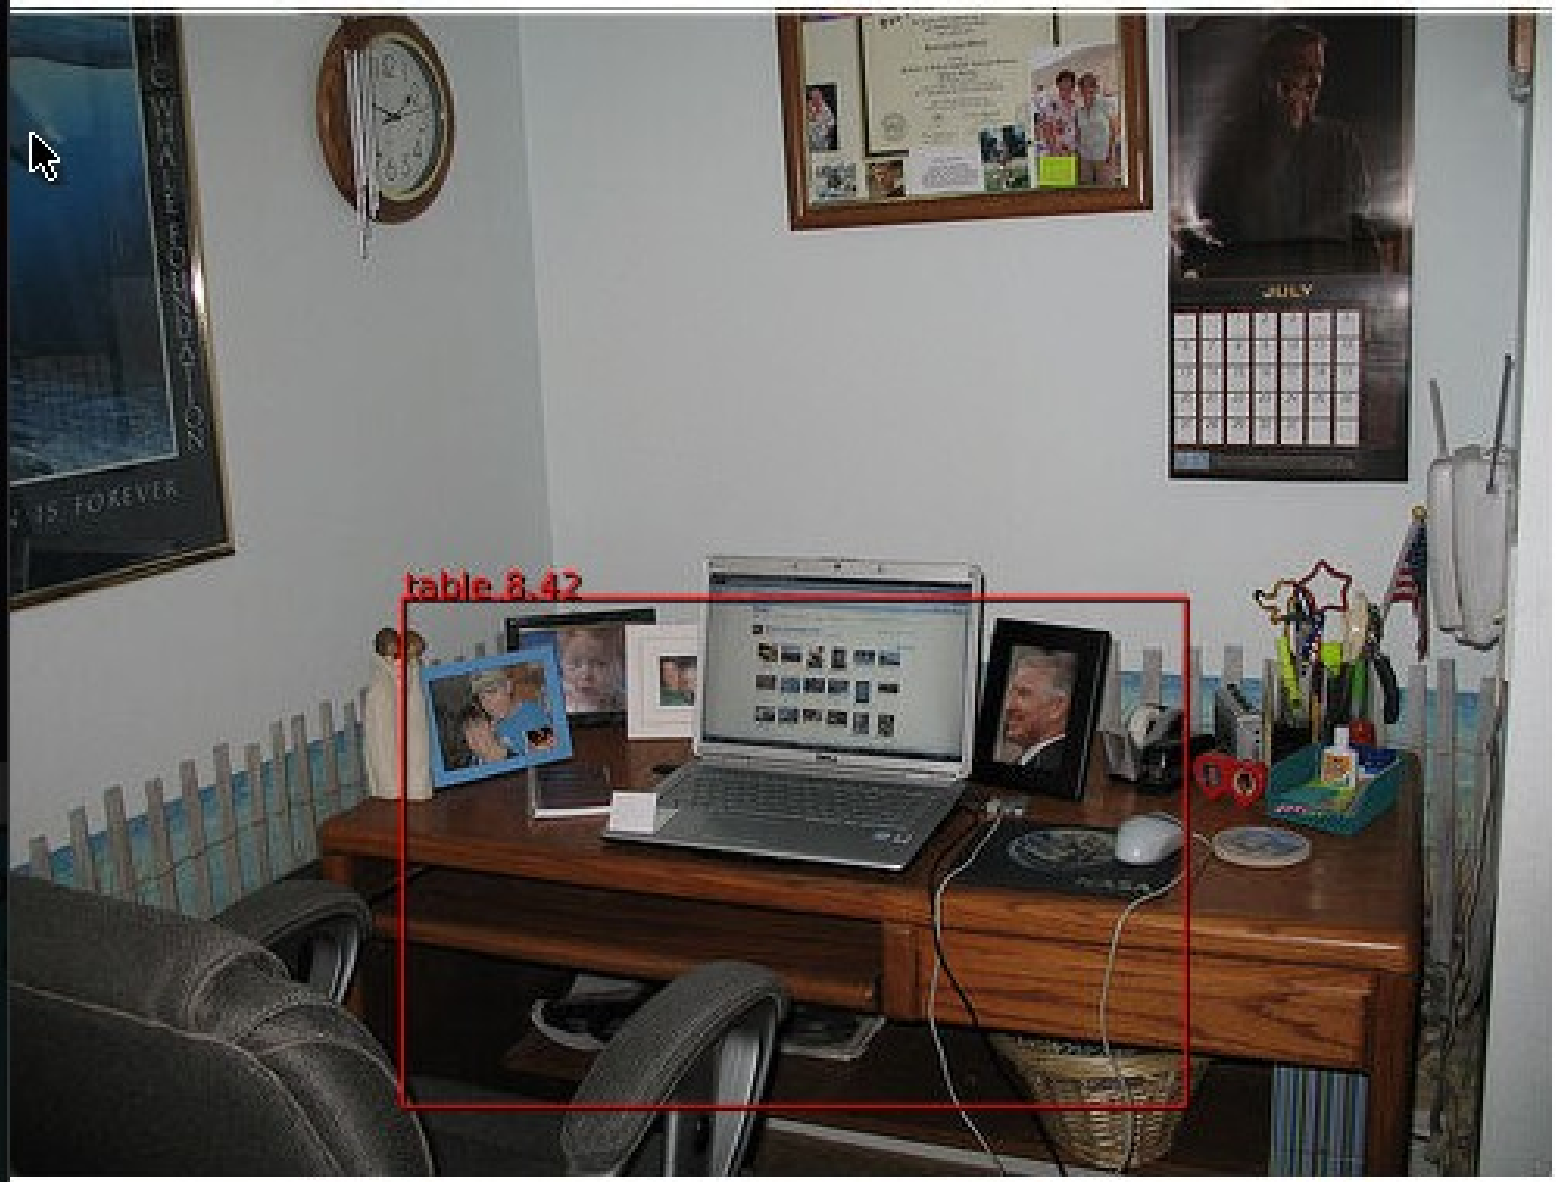

/home/snwiz/data/imagenet12/original/det/ILSVRC2013_DET_val/ILSVRC2012_val_00006665.JPEG
table conf 8.416754

Detection Examples

Detection Examples

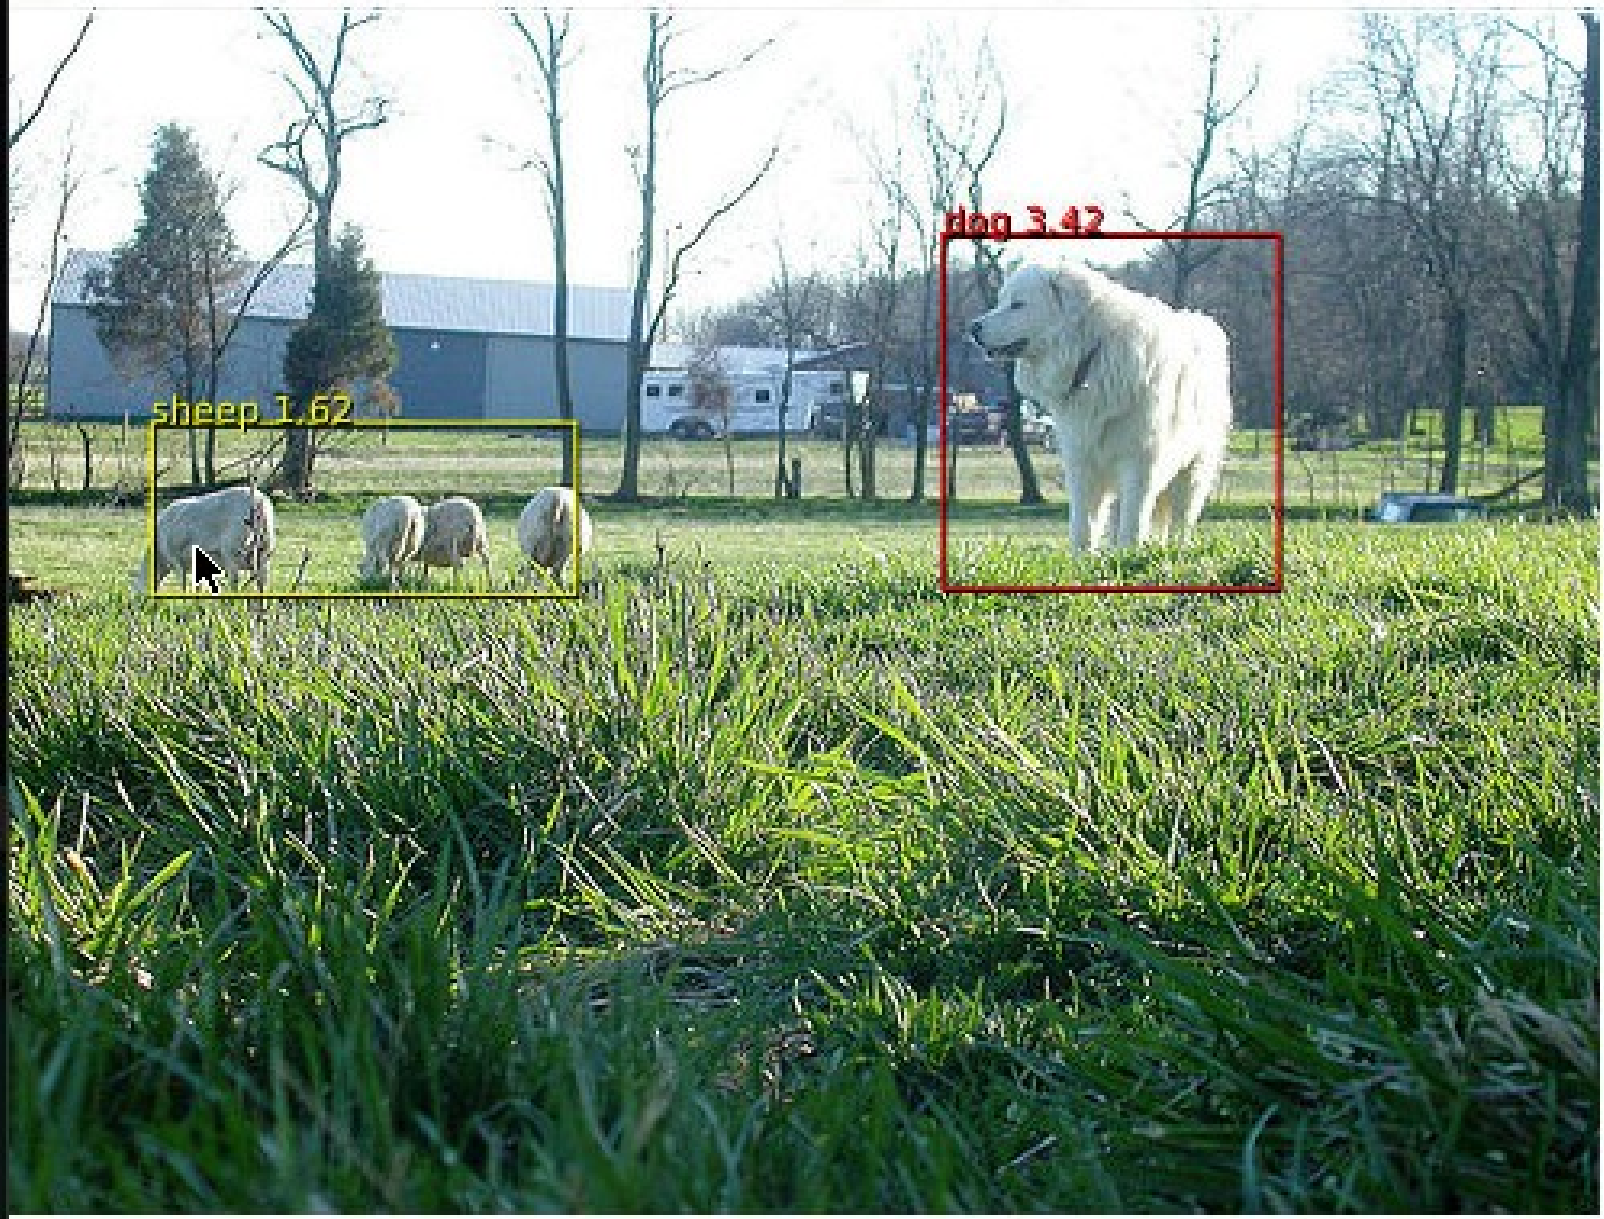

Detection Examples

/home/snwiz/data/imagenet12/original/det/ILSVRC2013_DET_test/ILSVRC2012_test_00091048.JPEG
person conf 17.898635
bow conf 15.628116

Detection Examples

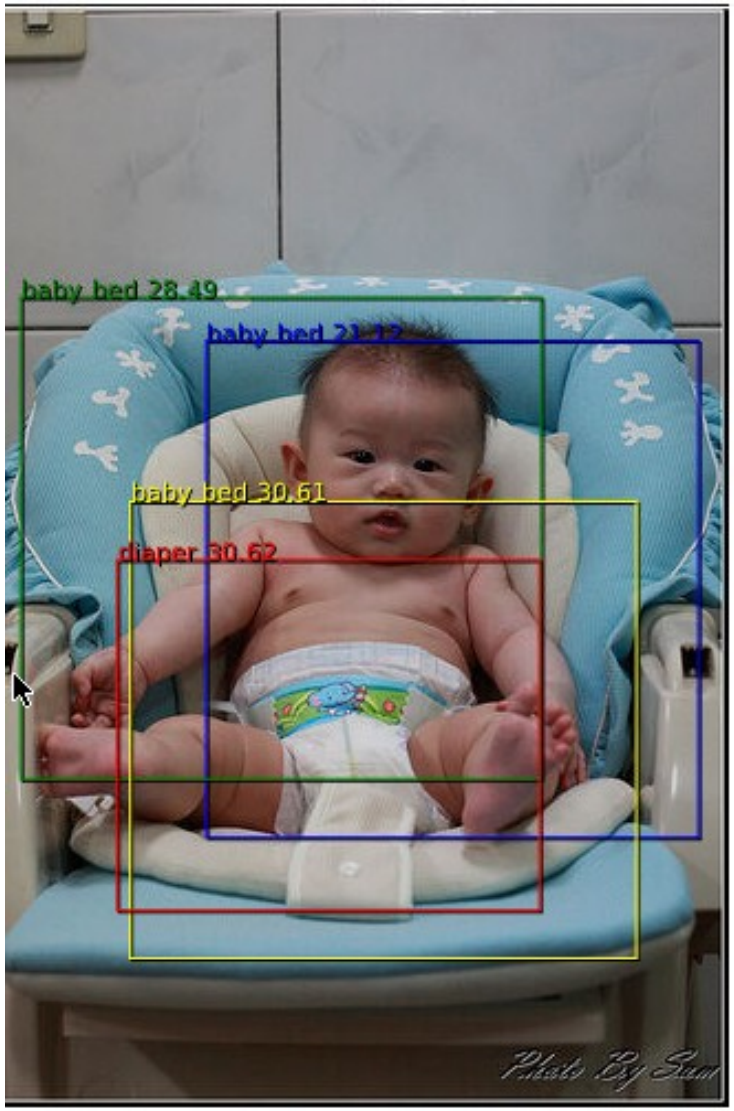

/home/snwiz/data/imagenet12/original/det/ILSVRC2013_DET_test/I
diaper conf 30.616426
baby bed conf 30.607744
baby bed conf 28.493934
baby bed conf 21.117424

Form

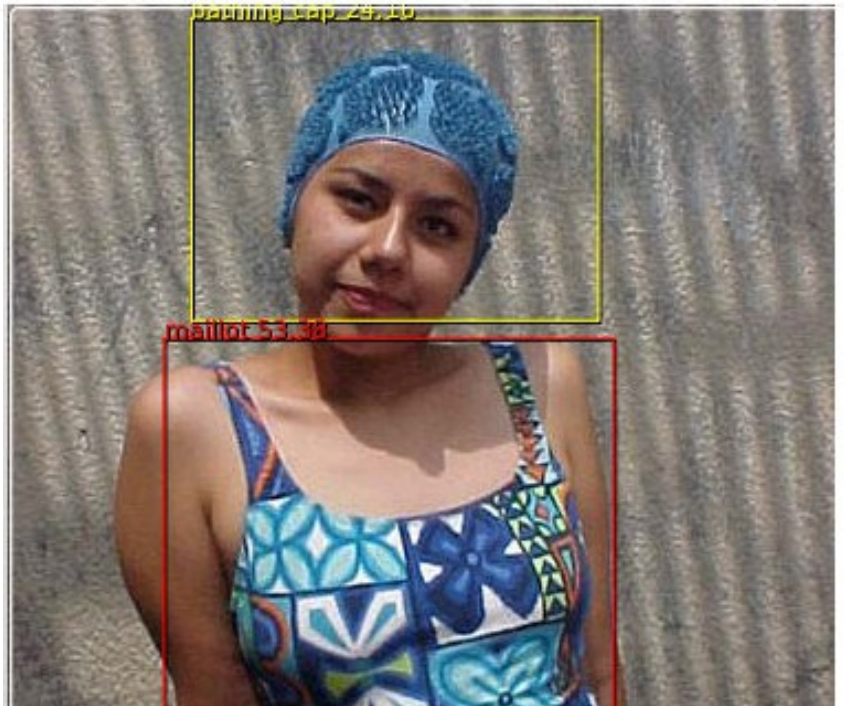

Form

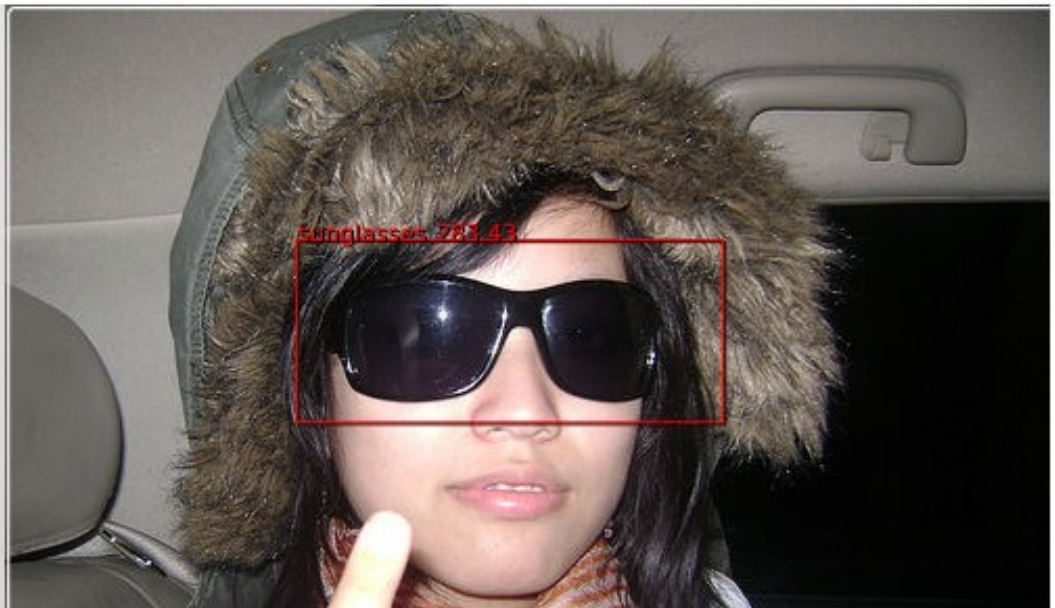

Detection Examples

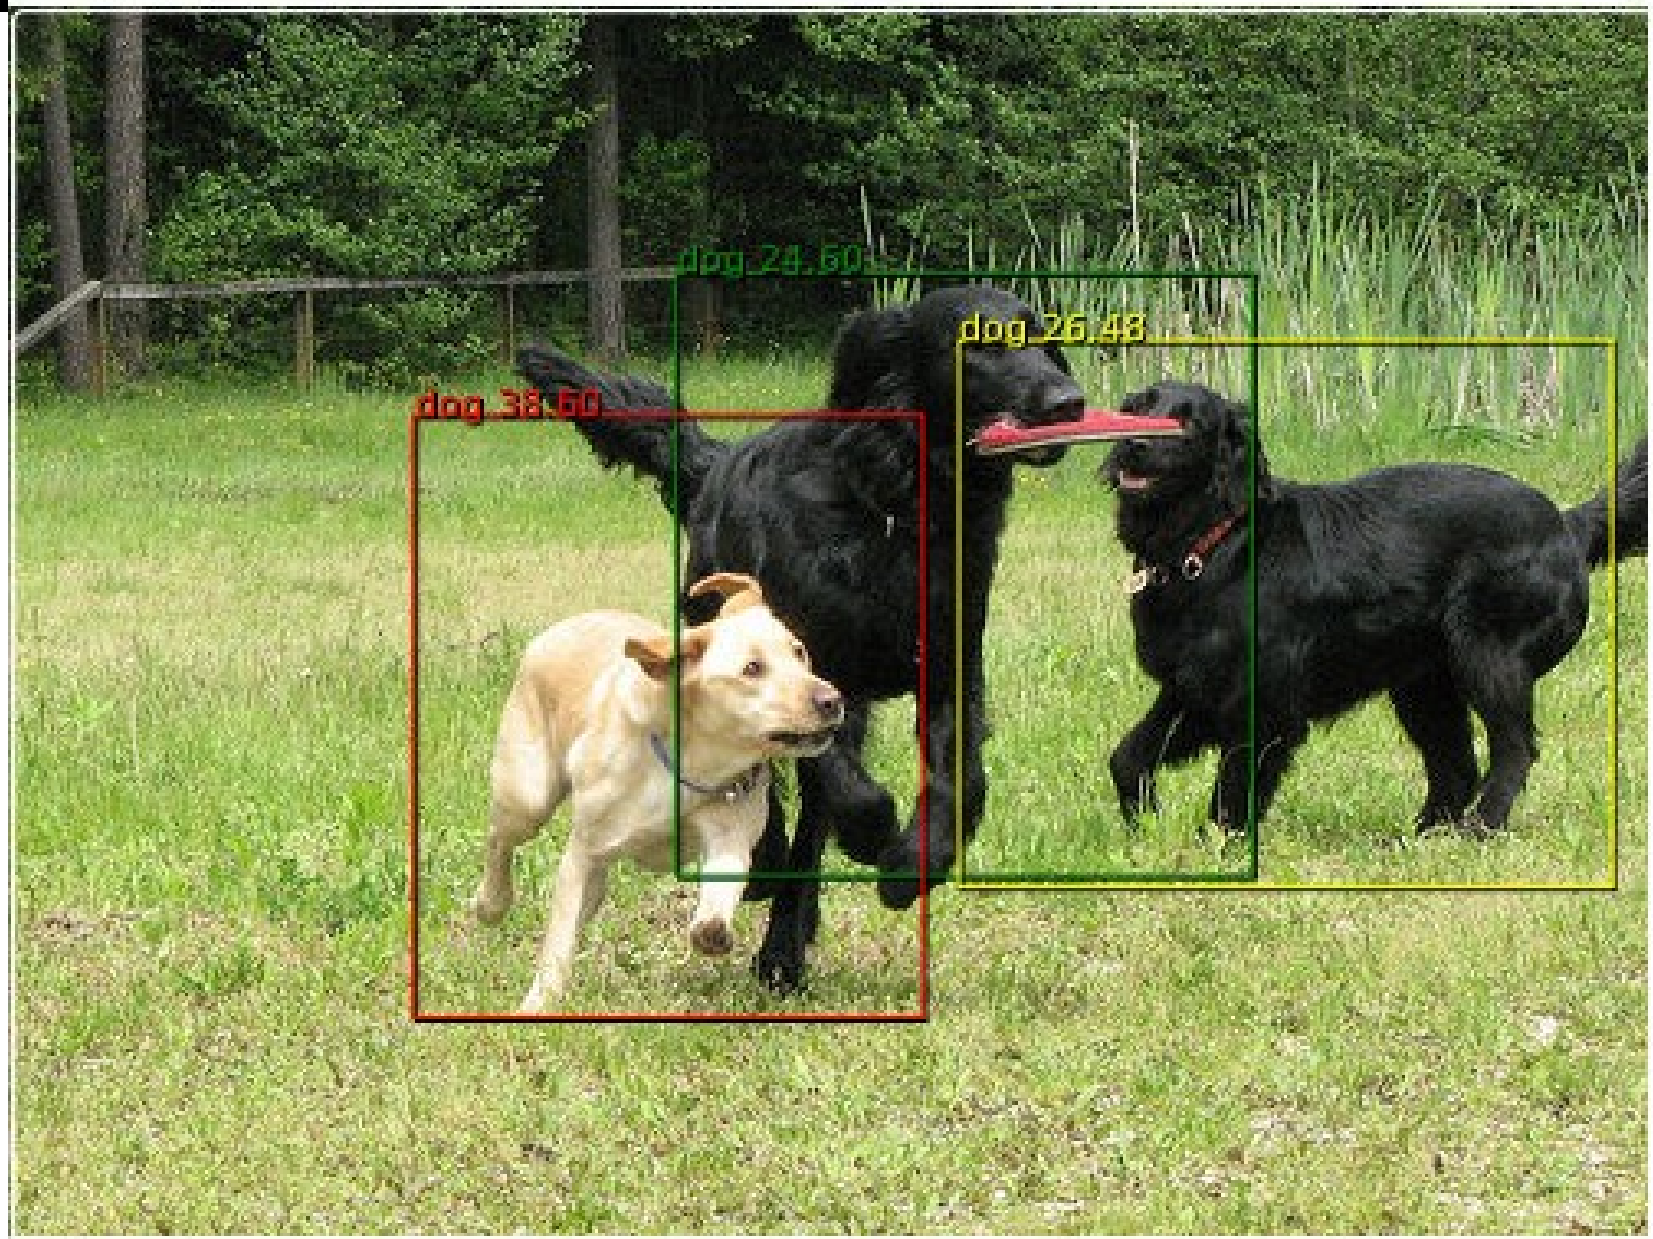

/home/snwiz/data/imagenet12/original/det/ILSVRC2013_DET_test/ILSVRC2012_test_00000172.JPEG
dog conf 38.603936

WE PROVIDI
THIS SERVICE
AT YOUR
person 9.06
french_horn 21.89

Detection: Difficult Examples

- Groundtruth is sometimes ambiguous or incomplete

Top predictions:
microwave (confidence 5.6)
refrigerator (confidence 2.5)

Groundtruth:
bowl
microwave

Top predictions:
person (confidence 6.0)

ILSVRC2012_val_00001273.JPEG

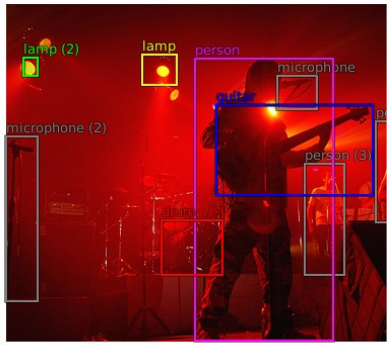

Groundtruth:
drum
lamp
lamp (2)
guitar
person
person (2)
person (3)
microphone
microphone (2)
microphone (3)

Top predictions:
artichoke (confidence 162.8)

ILSVRC2012_val_00001549.JPEG

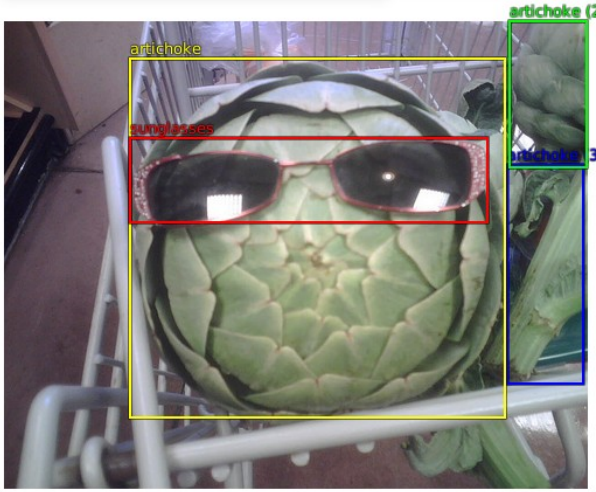

Groundtruth:
sunglasses
artichoke
artichoke (2)
artichoke (3)

Detection: Difficult Examples

Y LeCun

- Non-max suppression makes us miss many objects
  - Person behind instrument
- A bit of contextual post-processing would fix many errors

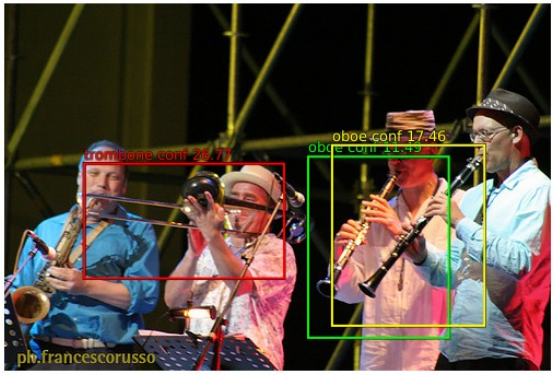

Top predictions:
trombone (confidence 26.8)
oboe (confidence 17.5)
oboe (confidence 11.5)

ILSVRC2012_val_00000614.JPEG

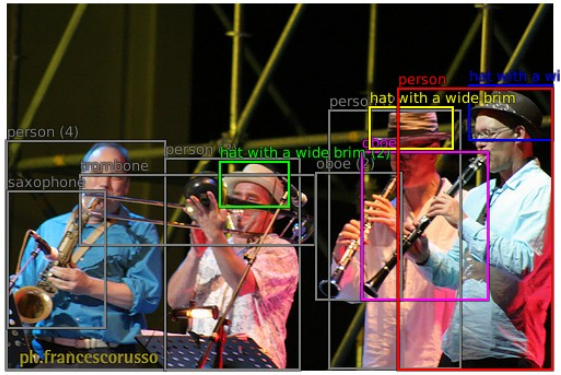

Groundtruth:
person
hat with a wide brim
hat with a wide brim (2)
hat with a wide brim (3)
oboe
oboe (2)
saxophone
trombone
person (2)
person (3)
person (4)

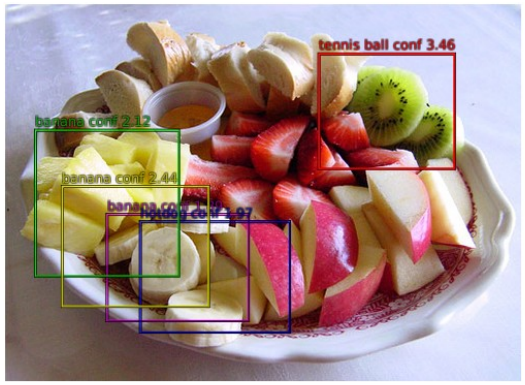

Top predictions:
tennis ball (confidence 3.5)
banana (confidence 2.4)
banana (confidence 2.1)
hotdog (confidence 2.0)
banana (confidence 1.9)

ILSVRC2012_val_00000320.JPEG

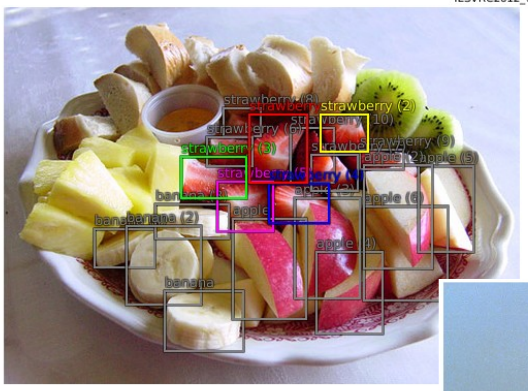

Groundtruth:
strawberry
strawberry (2)
strawberry (3)
strawberry (4)
strawberry (5)
strawberry (6)
strawberry (7)
strawberry (8)
strawberry (9)
strawberry (10)
apple
apple (2)
apple (3)

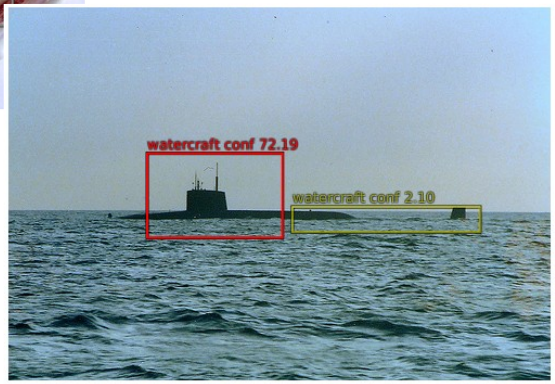

Top predictions:
watercraft (confidence 72.2)
watercraft (confidence 2.1)

ILSVRC2012_val_00000623.JPEG

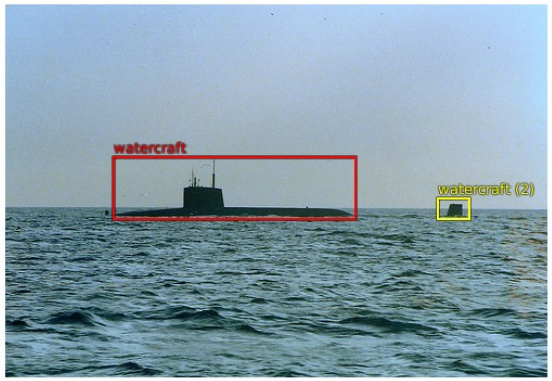

Groundtruth:
watercraft
watercraft (2)

Detection: Interesting Failures

Top predictions:
corkscrew (confidence 38.1)

ILSVRC2012_val_00000324.JPEG

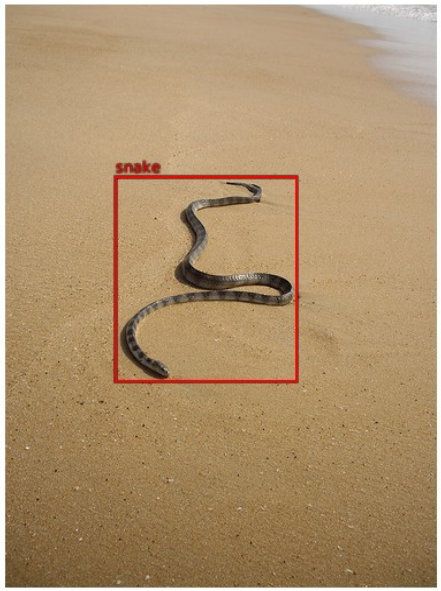

Groundtruth:
snake

Snake → Corkscrew

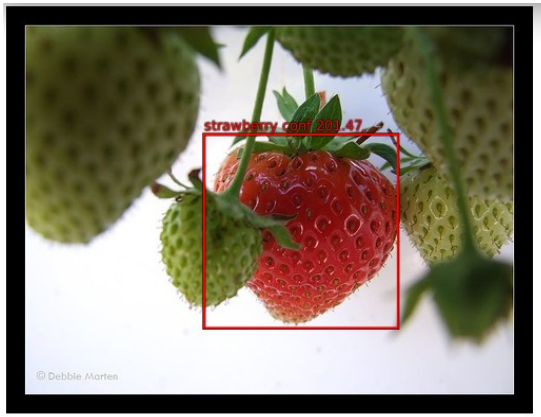

Top predictions:
strawberry (confidence 201.5)

ILSVRC2012_val_00000099.JPEG

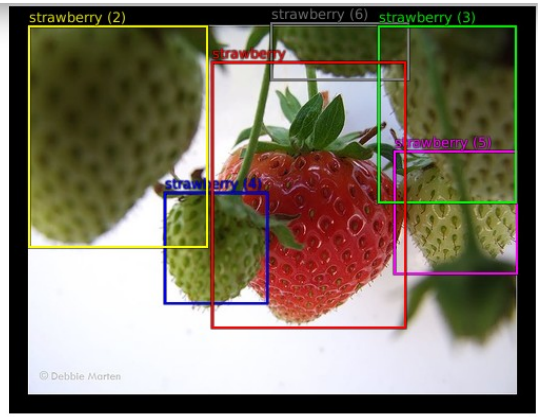

Groundtruth:
strawberry
strawberry (2)
strawberry (3)
strawberry (4)
strawberry (5)

Top predictions:
remote control (confidence 31.8)
filing cabinet (confidence 2.2)

ILSVRC2012_val_00000331.JPEG

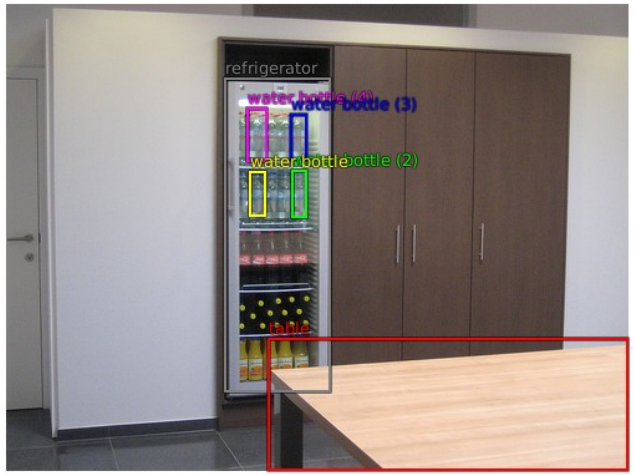

Groundtruth:
table
water bottle
water bottle (2)
water bottle (3)
water bottle (4)
refrigerator

Detection: Bad Groundtruth

One of the labelers likes ticks.....

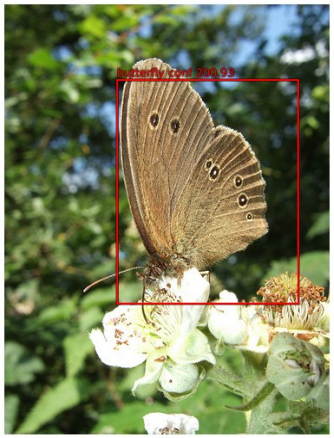

Top predictions:
butterfly (confidence 200.9)
ILSVRC2012_val_00035074.JPEG

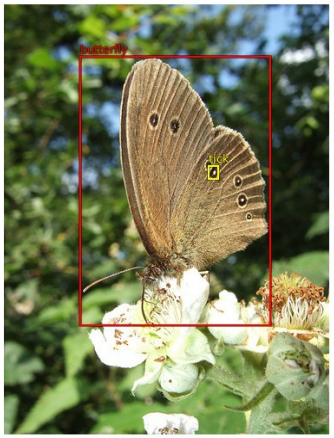

Groundtruth:
butterfly
tick

Top predictions:
butterfly (confidence 15.8)
ILSVRC2012_val_00012764.JPEG

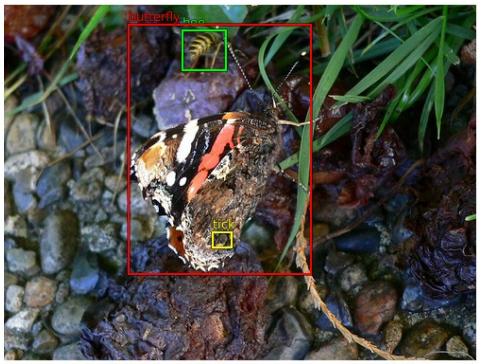

Groundtruth:
butterfly
tick
bee

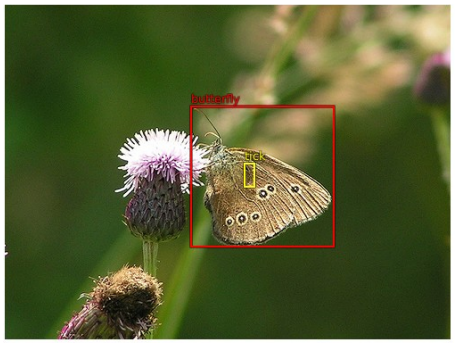

Groundtruth:
butterfly
tick

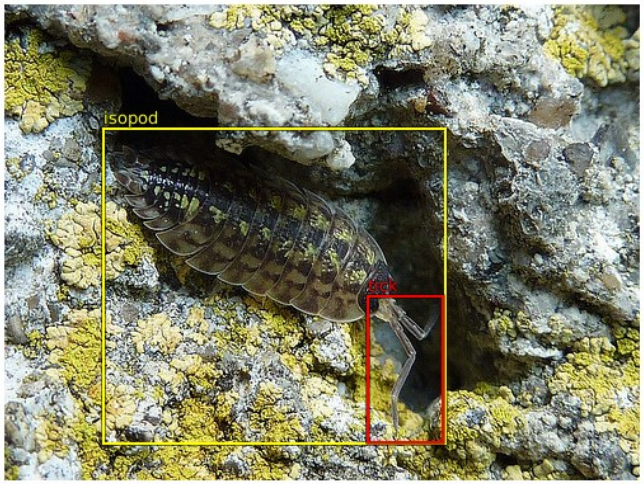

Groundtruth:
tick
isopod

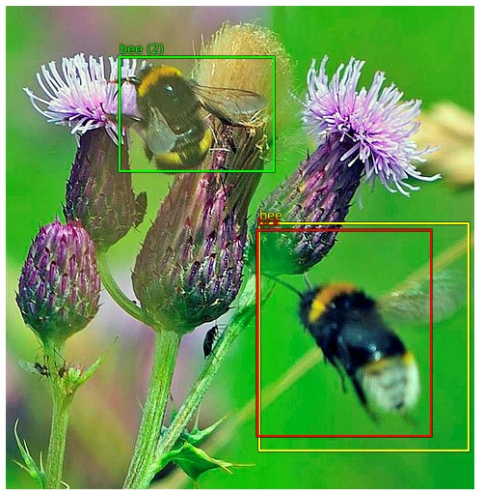

Groundtruth:
tick
bee
bee (2)

Top predictions:
snail (confidence 33.8)
ILSVRC2012_val_00023206.JPEG

Groundtruth:
snail
tick

ConvNets
As Generic
Feature Extractors

Cats vs Dogs

- Kaggle competition: Dog vs Cats

- Won by Pierre Sermanet (NYU):
- ImageNet network (OverFeat) last layers retrained on cats and dogs

Dogs vs. Cats

Finished

Wednesday, September 25, 2013 — Swag • 215 teams — Saturday, February 1, 2014

Dashboard

Leaderboard - Dogs vs. Cats

This competition has completed. This leaderboard reflects the final standings.

See someone using multiple accounts? Let us know.

#	Δ1w	Team Name * in the money	Score	Entries	Last Submission UTC (Best – Last Submission)
1	-	Pierre Sermanet *	0.98914	5	Sat, 01 Feb 2014 21:43:19 (-1.2h)
2	↑26	orchid *	0.98309	17	Sat, 01 Feb 2014 23:52:30
3	-	Owen	0.98171	15	Sat, 01 Feb 2014 17:04:40 (-1.3h)
4	new	Paul Covington	0.98171	3	Sat, 01 Feb 2014 23:05:20
5	↓3	Maxim Milakov	0.98137	24	Sat, 01 Feb 2014 18:20:58

OverFeat Features ->Trained Classifier on other datasets

Y LeCun

A. S. Razavian , H. Azizpour , J. Sullivan , S. Carlsson "CNN features off-the-shelf: An astounding baseline for recogniton", CVPR 2014, DeepVision Workshop.

http://www.csc.kth.se/cvap/cvg/DL/ots/

Comparing Best State of the Art Methods with Deep Represenations																			
	VOC07c	VOC12c	VOC12a	MIT67	SUN397	VOC07d	VOC10d	VOC11s	200Birds	102Flowers	H3Datt	UIUCatt	LFW	YTF	Paris6k	Oxford5k	Sculp6k	Holidays	UKB
best non-CNN results	70.5	82.2	69.6	64.0	47.2	34.3	40.4	47.6	56.8	80.7	69.9	~90.0	96.3	89.4	78.2	81.7	45.4	82.2	89.3
off-the-shelf ImageNet Model	80.13[10] 77.2[1]	82.7[10] 79.0[6]	-	69.0[1]	40.9[4]	46.2[2] 46.1[11]	44.1[11]	-	61.8[1] 58.8[4]	86.8[1]	73.0[1]	91.5[1]	-	-	79.5[1]	68.0[1]	42.3[1]	84.3[1]	91.1[1]
off-the-shelf ImageNet Model + rep learning	-	-	-	68.9[3]	52.0[3]	-	-	-	65.0[4]	-	-	-	-	-	-	-	-	80.2[3]	-
fine-tuned ImageNet Model	82.42[10] 77.7[5]	83.2[10] 82.8[5]	70.2[5]	-	-	58.5[2]	53.7[2]	47.9[2]	-	-	-	-	-	-	-	-	-	-	-
Other Deep Learning Models	-	-	-	-	-	-	-	-	-	-	79.0[7]	-	97.35[8]	91.4[8]	-	-	-	-	-

VOC07c:	Pascal VOC 2007 (Object Image Classification)
VOC12c:	Pascal VOC 2012 (Object Image Classification)
VOC12a:	Pascal VOC 2012 (Action Image Classification)
MIT67:	MIT 67 Indoor Scenes(Scene Image Classification)
VOC07d:	PASCAL VOC 2007 (Object Detection)
VOC10d:	PASCAL VOC 2010 (Object Detection)
VOC12d:	PASCAL VOC 2012 (Object Detection)
VOC11s:	PASCAL VOC 2011 (Object Category Segmentatic
200Birds:	UCSD-Caltech 2011-200 Birds dataset (Fine-grained Recognition)
102Flowers:	Oxford 102 Flowers (Fine-grained Recognition)
H3Datt:	H3D poselets Human 9 Attributes (Attribute Detection)
UIUCatt:	UIUC object attributes (Attribute Detection)
LFW:	Labelled Faces in the Wild (Metric Learning)
Oxford5k:	Oxford 5k Buildings Dataset (Instance Retrieval)
Paris6k:	Paris 6k Buildings Dataset (Instance Retrieval)
Sculp6k:	Oxford Sculptures Dataset (Instance Retrieval)
Holidays:	INRIA Holidays Scenes Dataset (Instance Retrieval)
UKB:	Uni. of Kentucky Retrieval Benchmark Dataset (Instance Retrieval)

OverFeat Features + Classifier on various datasets

	Dataset	Performance	Score
[Sermanet et al 2014]: OverFeat (fine-tuned features for each task) (tasks are ordered by increasing difficulty)			
• image classification	ImageNet LSVRC 2013	competitive	13.6 % error
	Dogs vs Cats Kaggle challenge 2014	**state of the art**	98.9%
• object localization	ImageNet LSVRC 2013	**state of the art**	29.9% error
• object detection	ImageNet LSVRC 2013	**state of the art**	24.3% mAP
[Razavian et al, 2014]: public OverFeat library (no retraining) + SVM **(simplest approach possible on purpose, no attempt at more complex classifiers)** (tasks are ordered by "distance" from classification task on which OverFeat was trained)			
• image classification	Pascal VOC 2007	competitive	73.9% mAP
• scene recognition	MIT-67	competitive	58.4% mAP
• fine grained recognition	Caltech-UCSD Birds 200-2011	competitive	53.3% mAP
	Oxford 102 Flowers	competitive	74.70% mAP
• attribute detection	UIUC 64 object attributes	**state of the art**	89.0% mAUC
	H3D Human Attributes	**state of the art**	70.78% mAP
• image retrieval (search by image similarity)	Oxford 5k buildings	?	0.52
	Paris 6k buildings	?	0.676
	Sculp6k	?	0.269
	Holidays	competitive	0.646
	UKBench	relatively poor	3.05

Pierre Sermanet, David Eigen, Xiang Zhang, Michael Mathieu, Rob Fergus, Yann LeCun, **OverFeat: Integrated Recognition, Localization and Detection using Convolutional Networks**, http://arxiv.org/abs/1312.6229, ICLR 2014
Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, Stefan Carlsson, **CNN Features off-the-shelf: an Astounding Baseline for Recognition**, http://arxiv.org/abs/1403.6382, DeepVision CVPR 2014 workshop

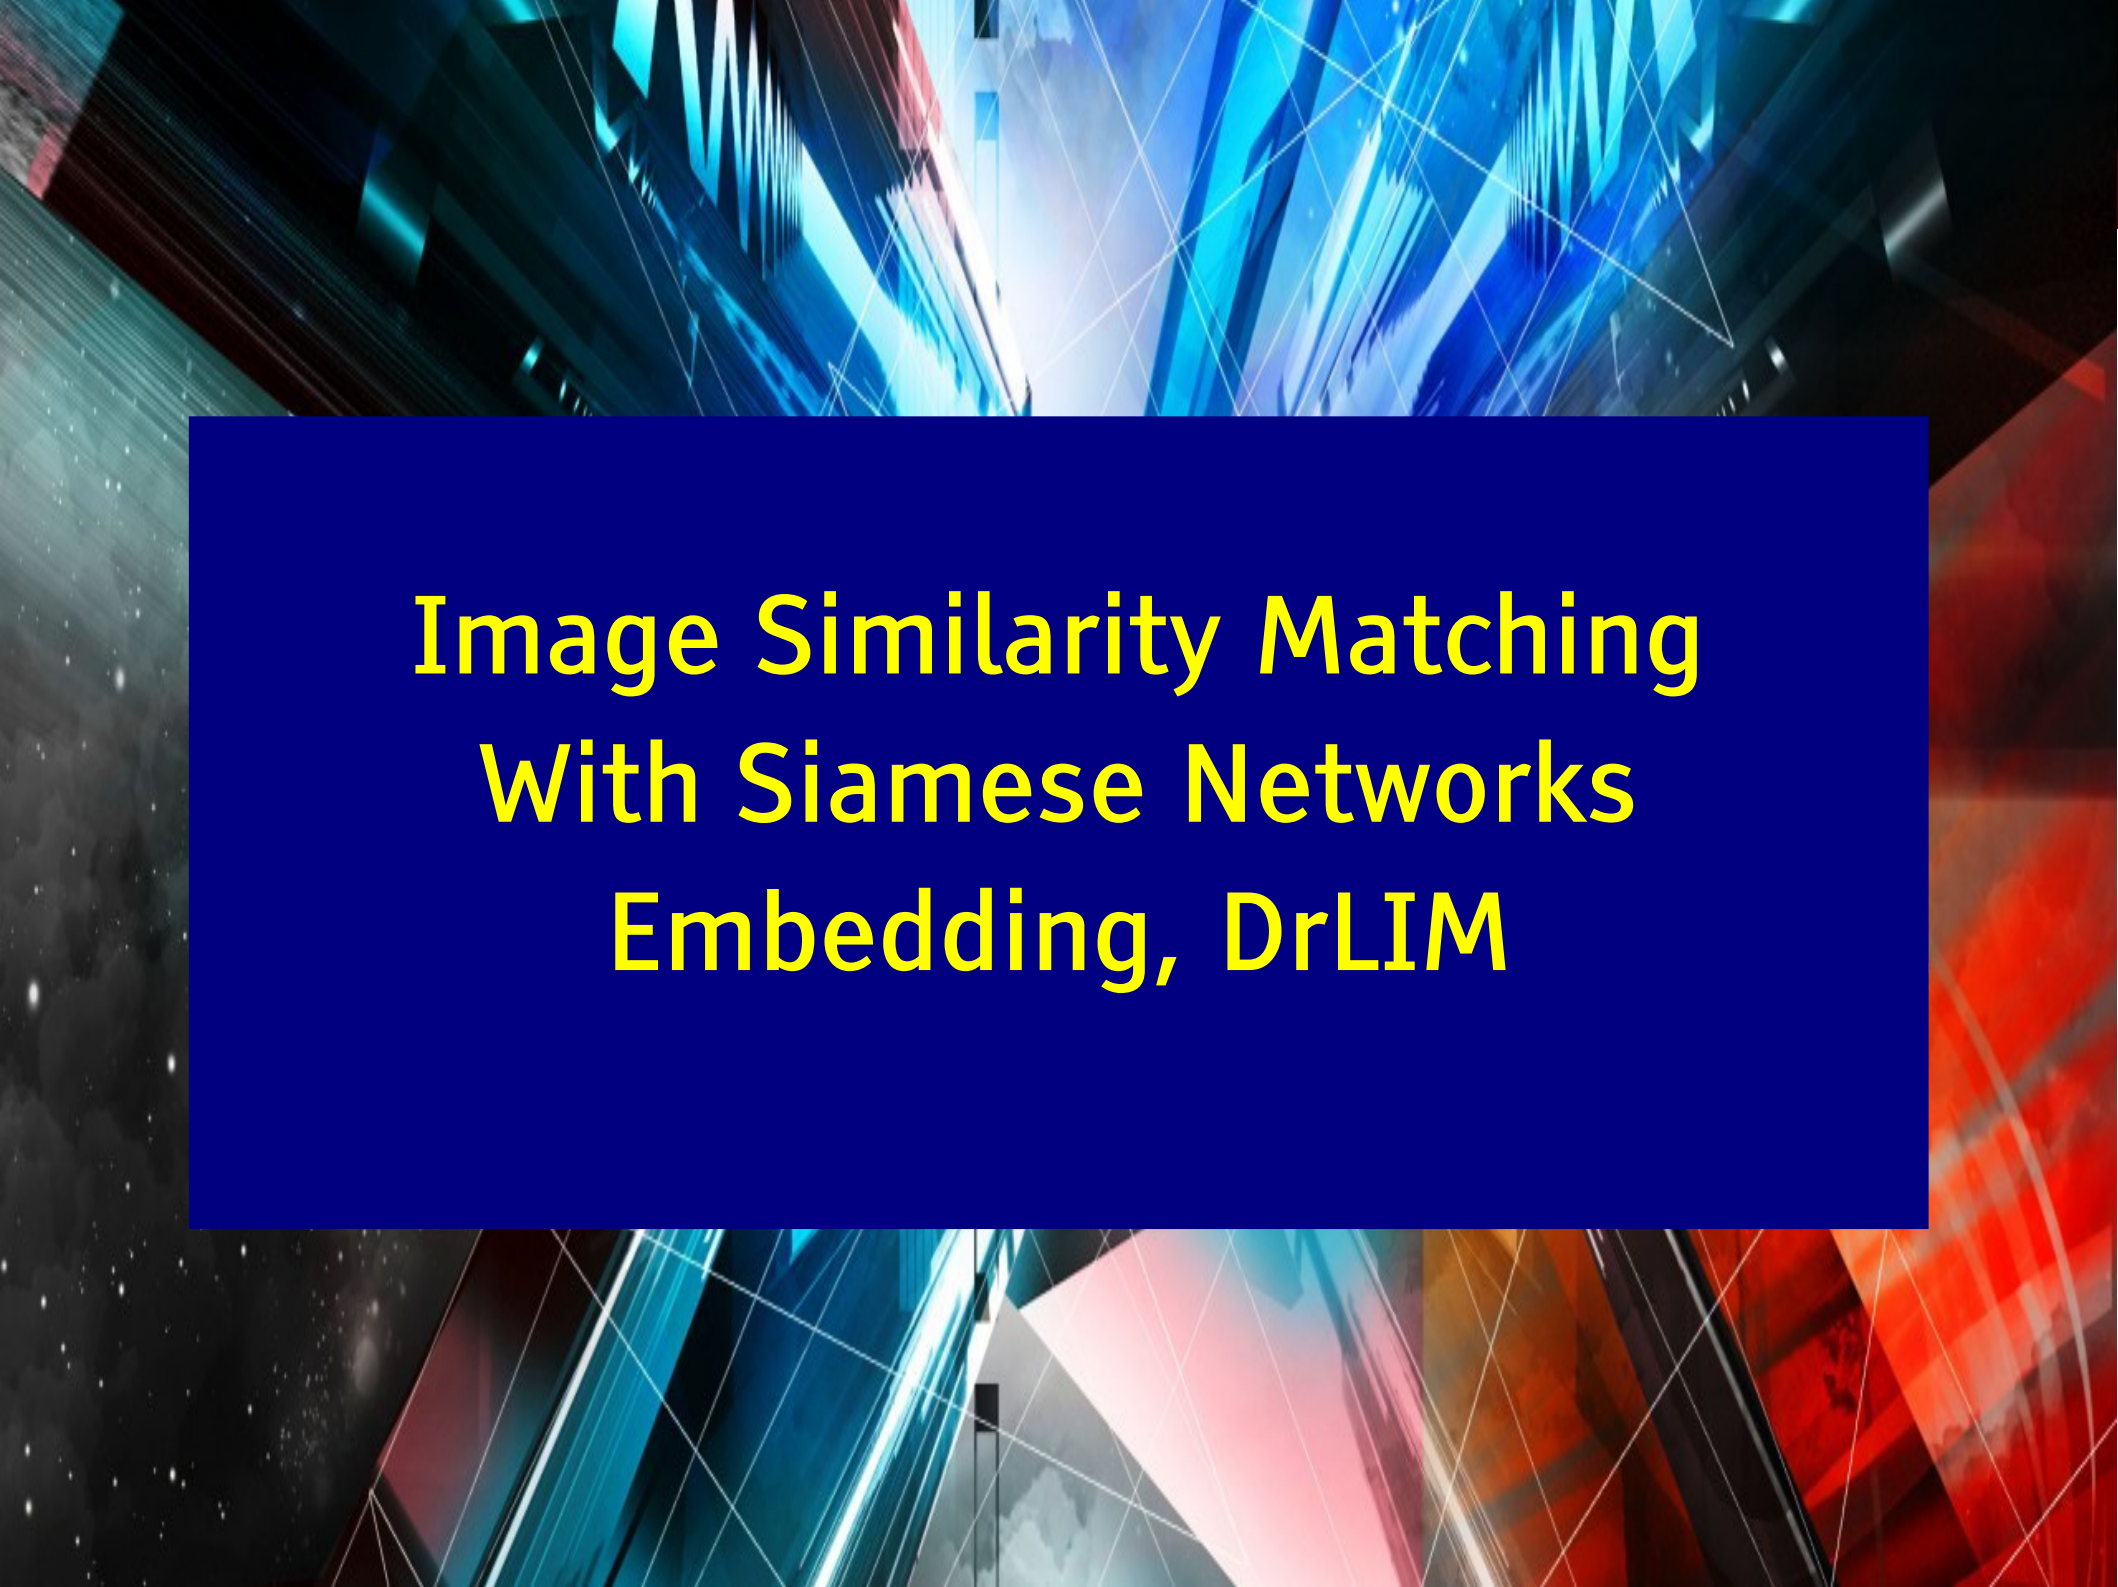
Image Similarity Matching
With Siamese Networks
Embedding, DrLIM

DrLIM: Metric Learning

Dimensionality Reduction by Learning an Invariant Mapping

- **Step 1**: Construct neighborhood graph.
- **Step 2:** Choose a parameterized family of functions.
- **Step 3:** Optimize the parameters such that:
  - Outputs for similar samples are pulled closer.
  - Outputs for dissimilar samples are pushed away.

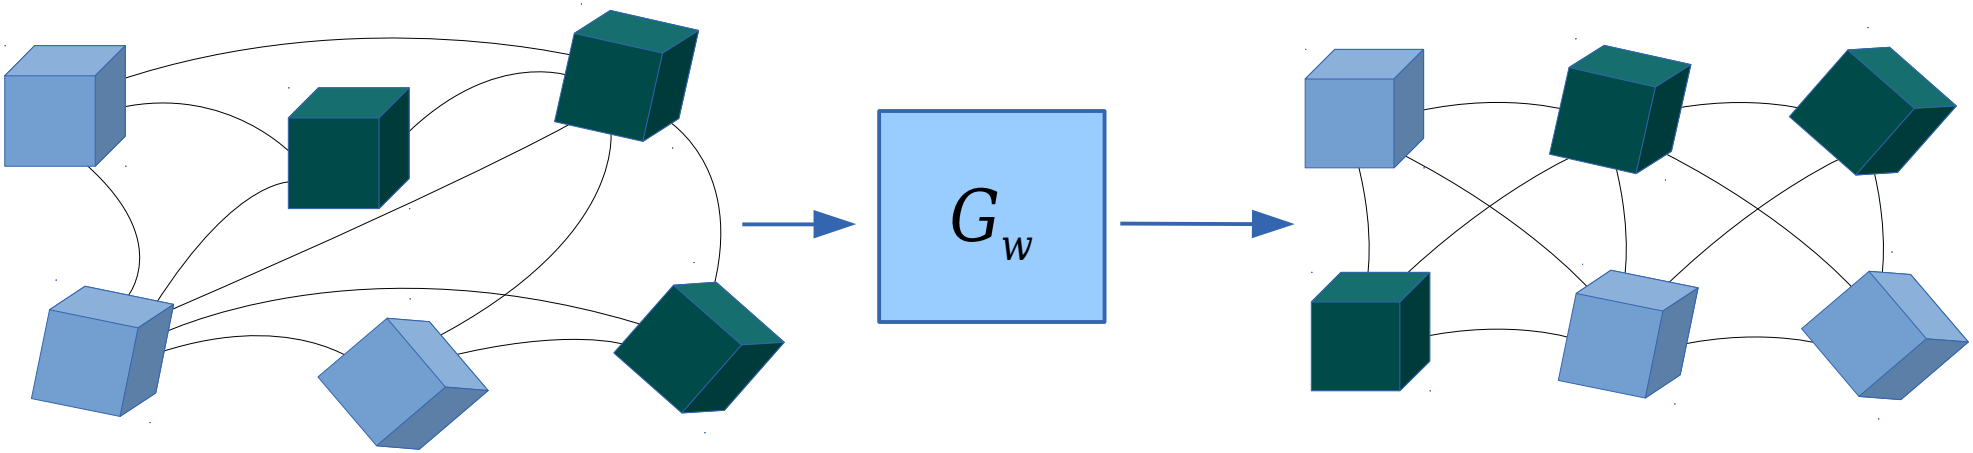

joint work with Sumit Chopra: Hadsell et al. CVPR 06; Chopra et al., CVPR 05

Siamese Architecture

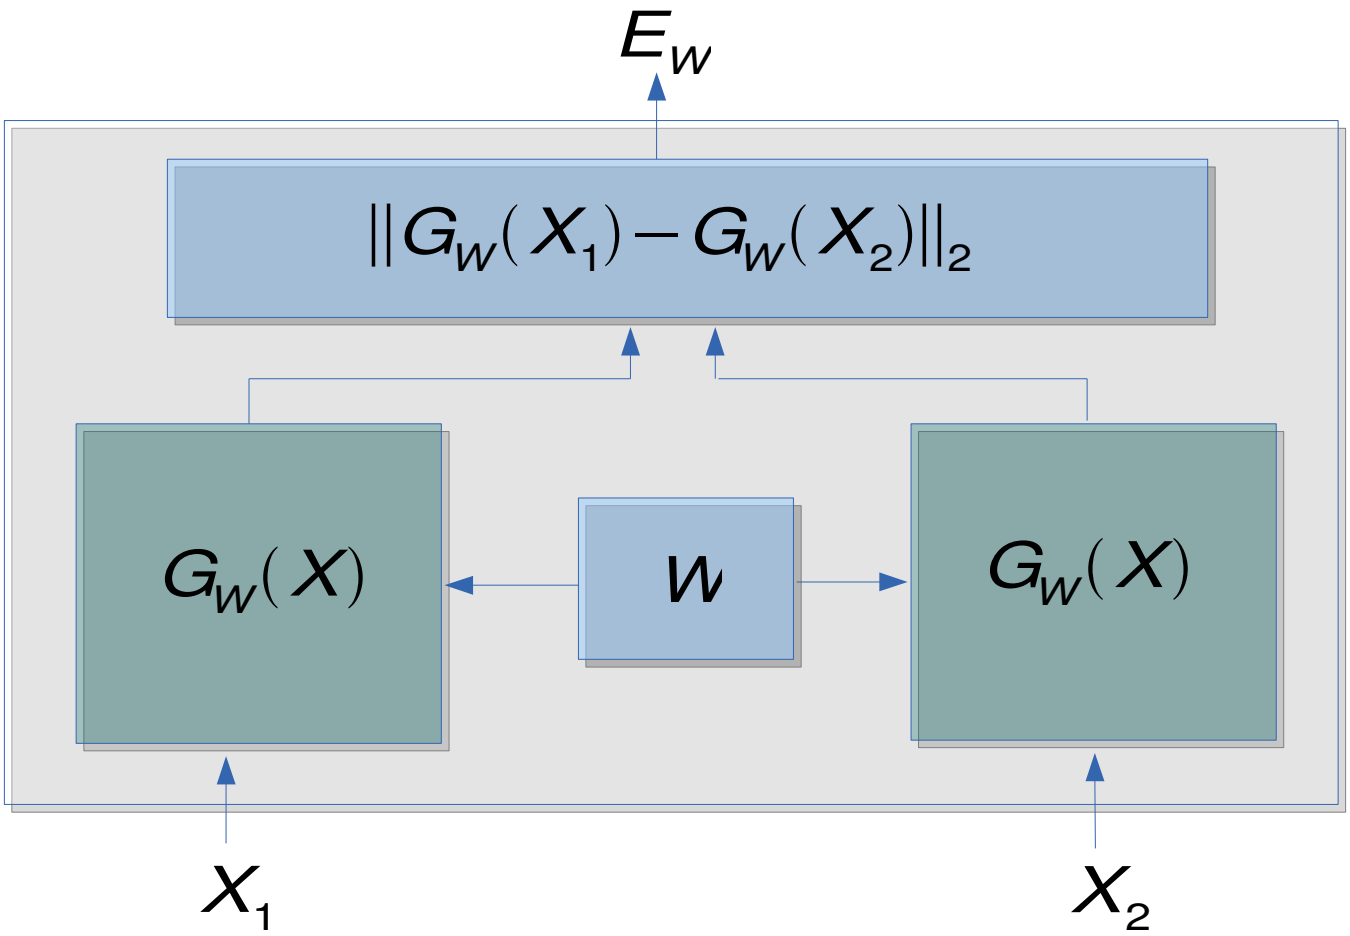

Siamese Architecture [Bromley, Sackinger, Shah, LeCun 1994]

Siamese Architecture and loss function

Loss function:

- Outputs corresponding to input samples that are neighbors in the neigborhood graph should be nearby
- Outputs for input samples that are not neighbors should be far away from each other

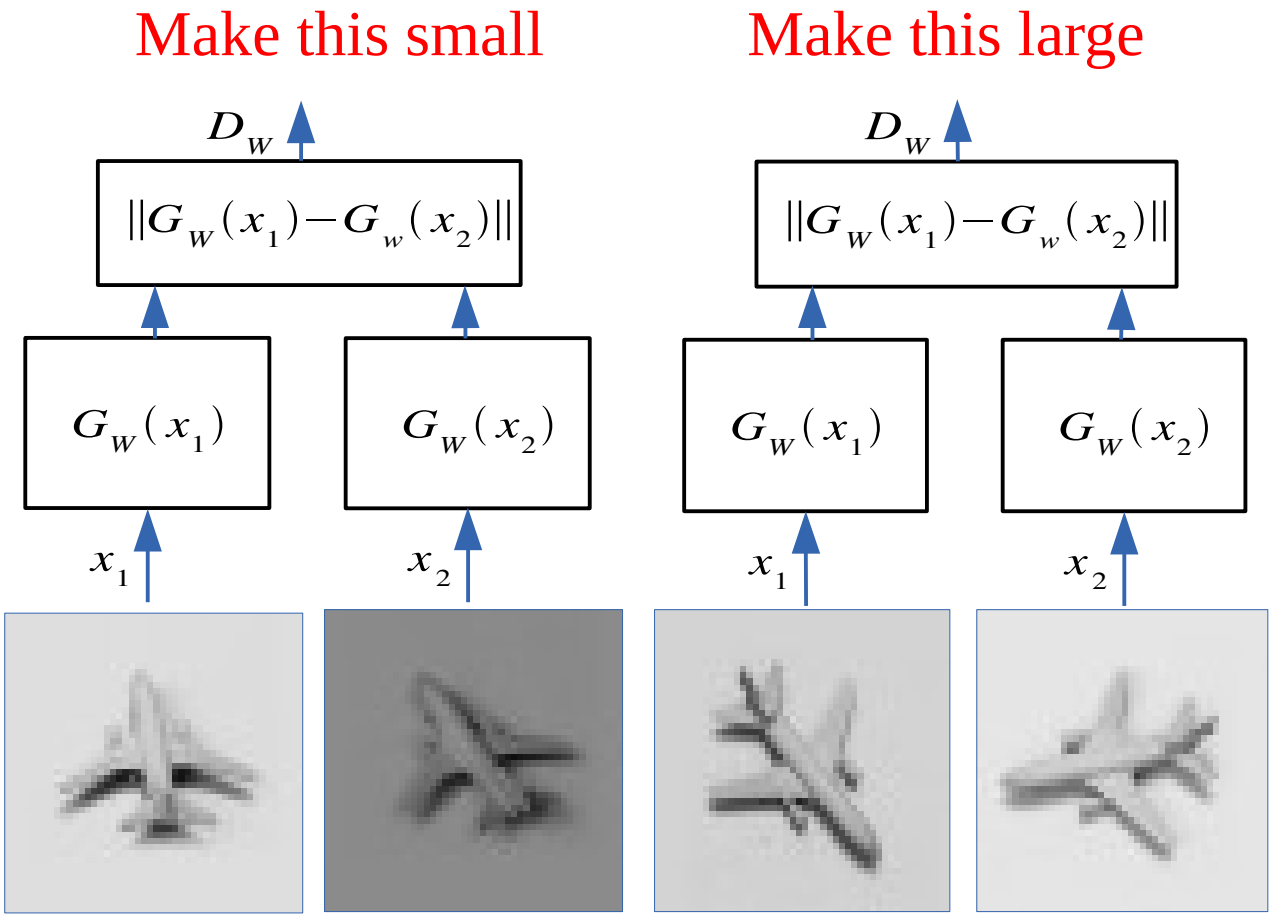

Face Recognition:DeepFace (Facebook AI Research)

- [Taigman et al. CVPR 2014]
  - Alignment
  - Convnet

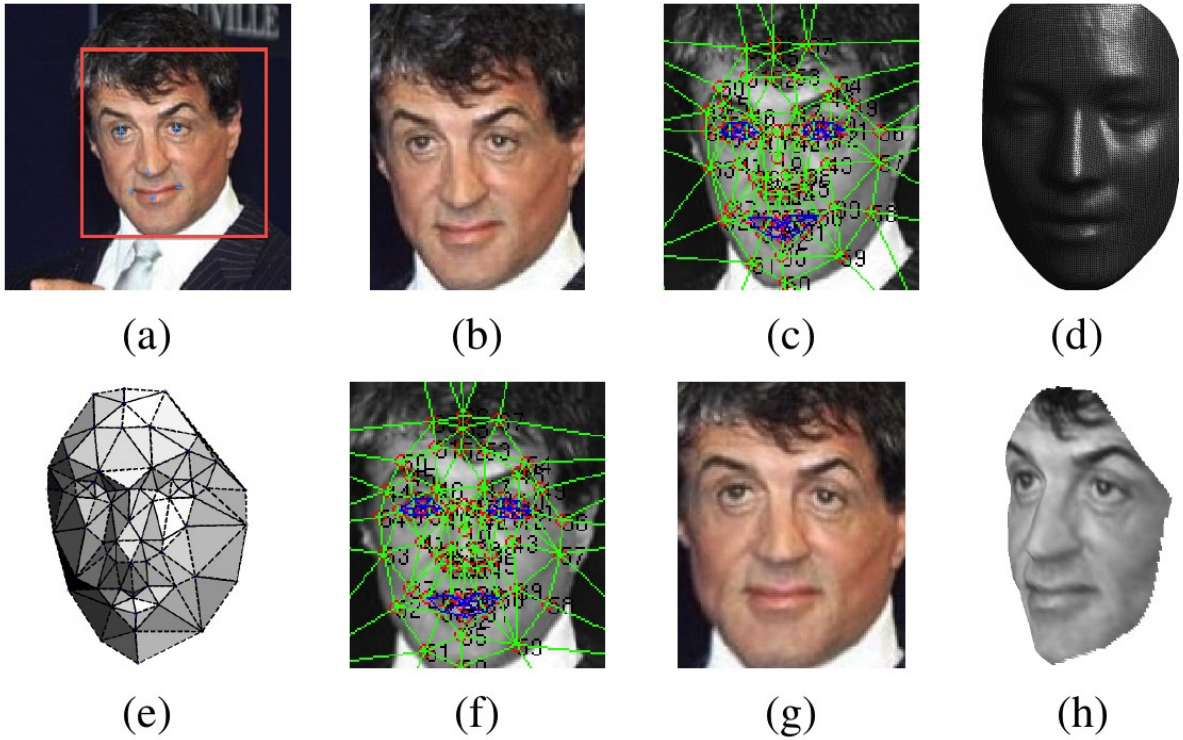

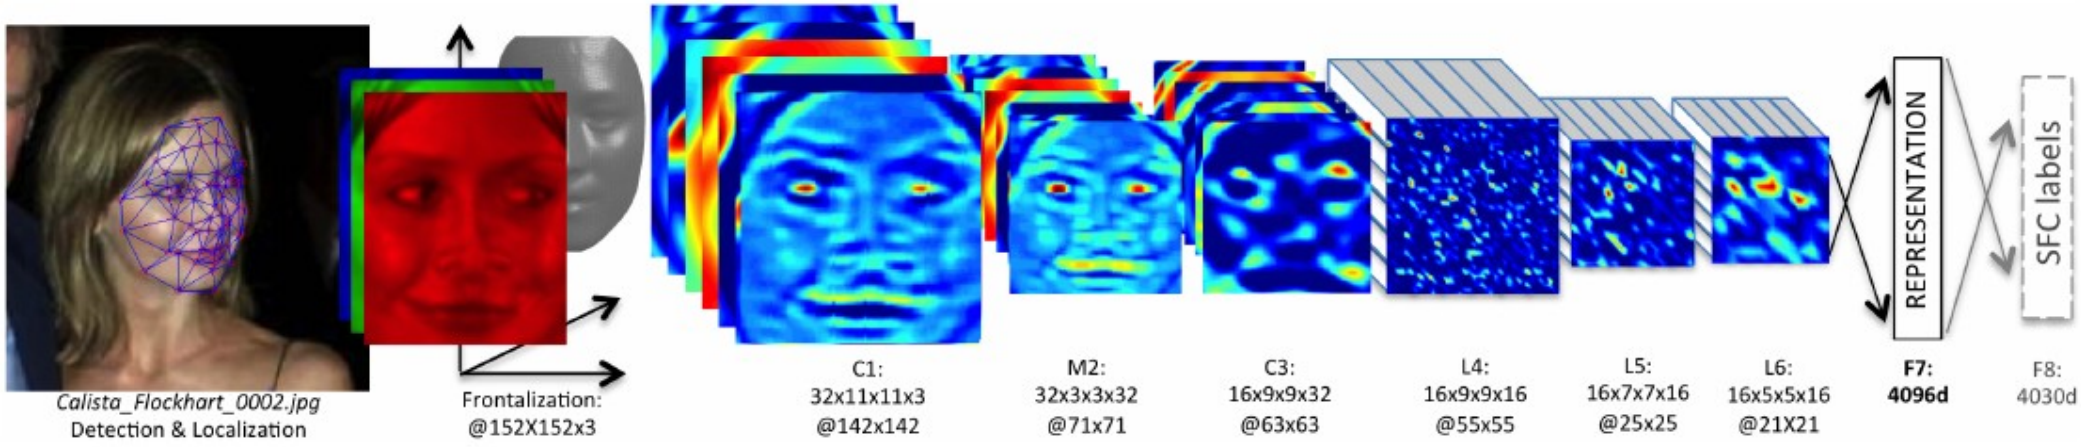

Face Recognition:DeepFace (Facebook AI Research)

Performance on Labeled Face in the Wild dataset (LFW)

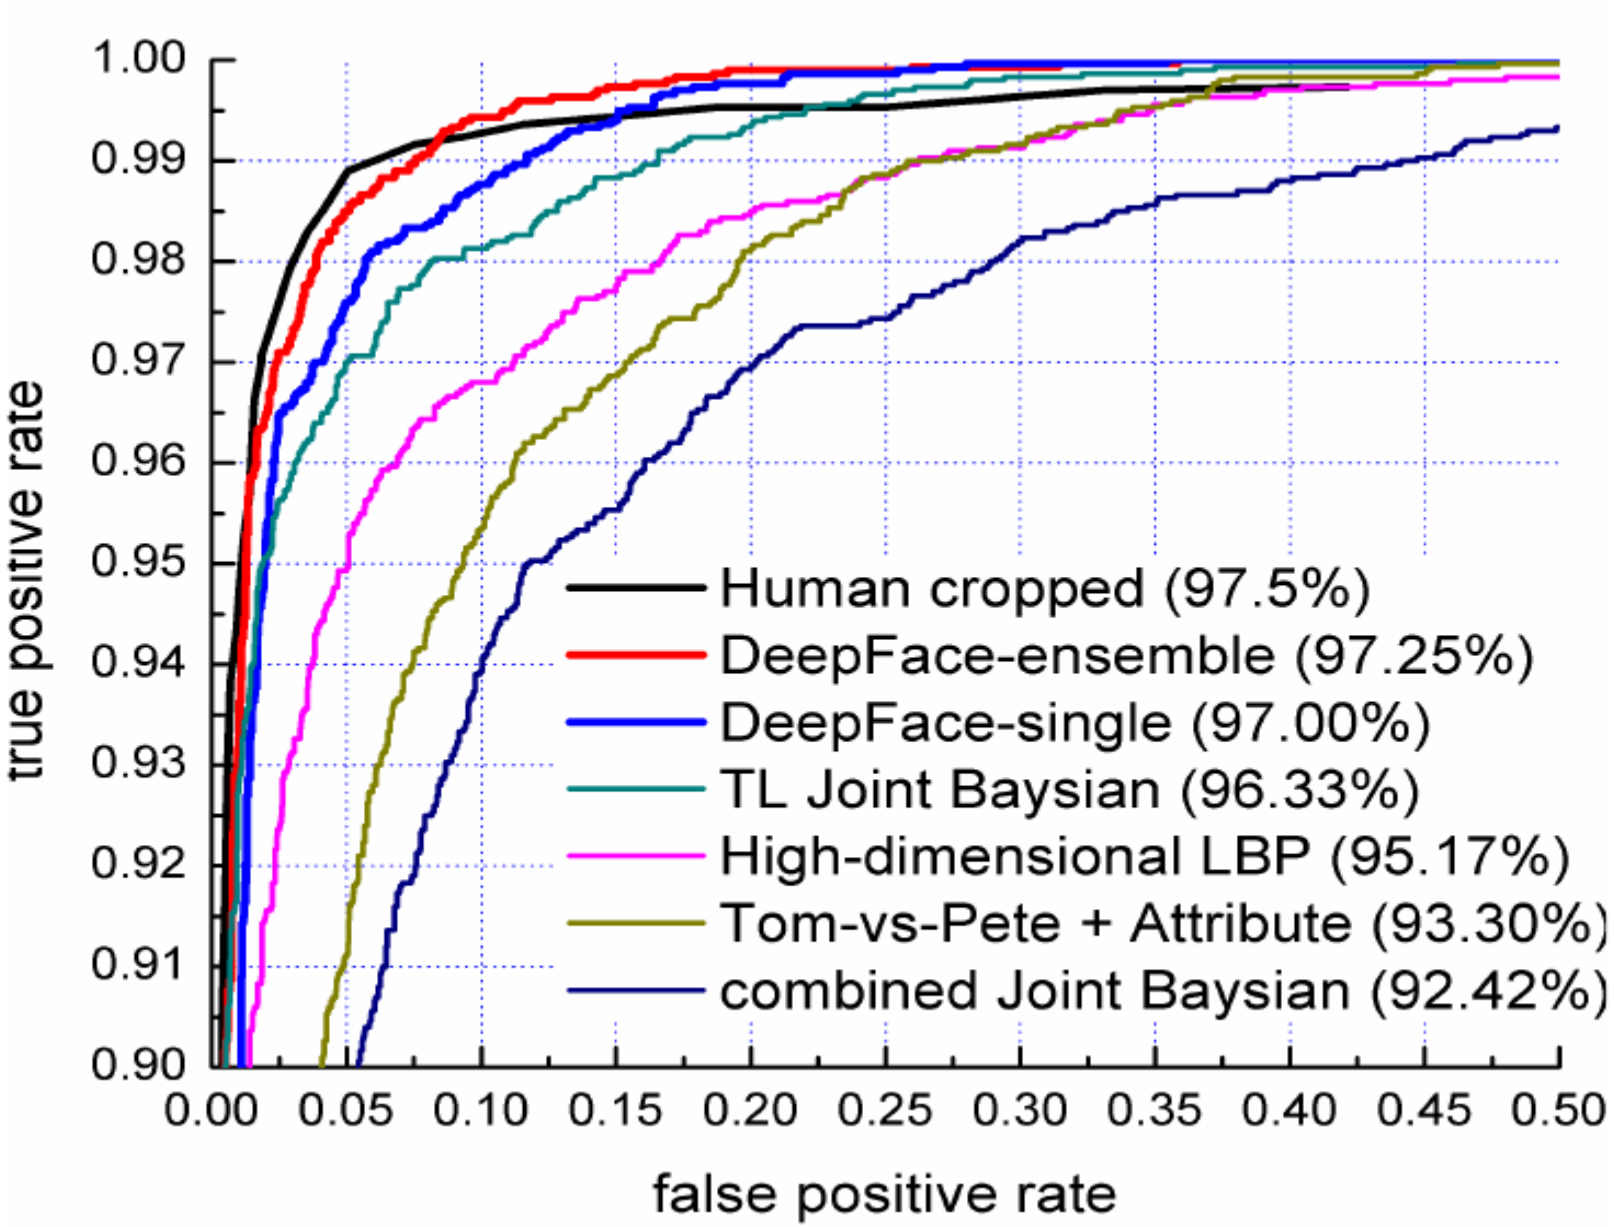

DeepFace: performance

Method	Accuracy	Protocol
Joint Bayesian [6]	0.9242 ±0.0108	restricted
Tom-vs-Pete [4]	0.9330 ±0.0128	restricted
High-dim LBP [7]	0.9517 ±0.0113	restricted
TL Joint Bayesian [5]	0.9633 ±0.0108	restricted
DeepFace-single	**0.9592** ±0.0092	unsupervised
DeepFace-single	**0.9700** ±0.0087	restricted
DeepFace-ensemble	**0.9715** ±0.0084	restricted
DeepFace-ensemble	**0.9725** ±0.0081	unrestricted
Human, cropped	0.9753	

Table 3. Comparison with the state-of-the-art on the *LFW* dataset.

Method	Accuracy (%)	AUC	EER
MBGS+SVM- [31]	78.9 ±1.9	86.9	21.2
APEM+FUSION [22]	79.1 ±1.5	86.6	21.4
STFRD+PMML [9]	79.5 ±2.5	88.6	19.9
VSOF+OSS [23]	79.7 ±1.8	89.4	20.0
DeepFace-single	**91.4** ±1.1	96.3	8.6

Table 4. Comparison with the state-of-the-art on the *YTF* dataset.

Network	Error (*SFC*)	Accuracy (*LFW*)
DeepFace-gradient	8.9%	0.9582 ±0.0118
DeepFace-align2D	9.5%	0.9430 ±0.0136
DeepFace-Siamese	NA	0.9617 ±0.0120

Table 2. The performance of various individual *DeepFace* networks and the Siamese network.

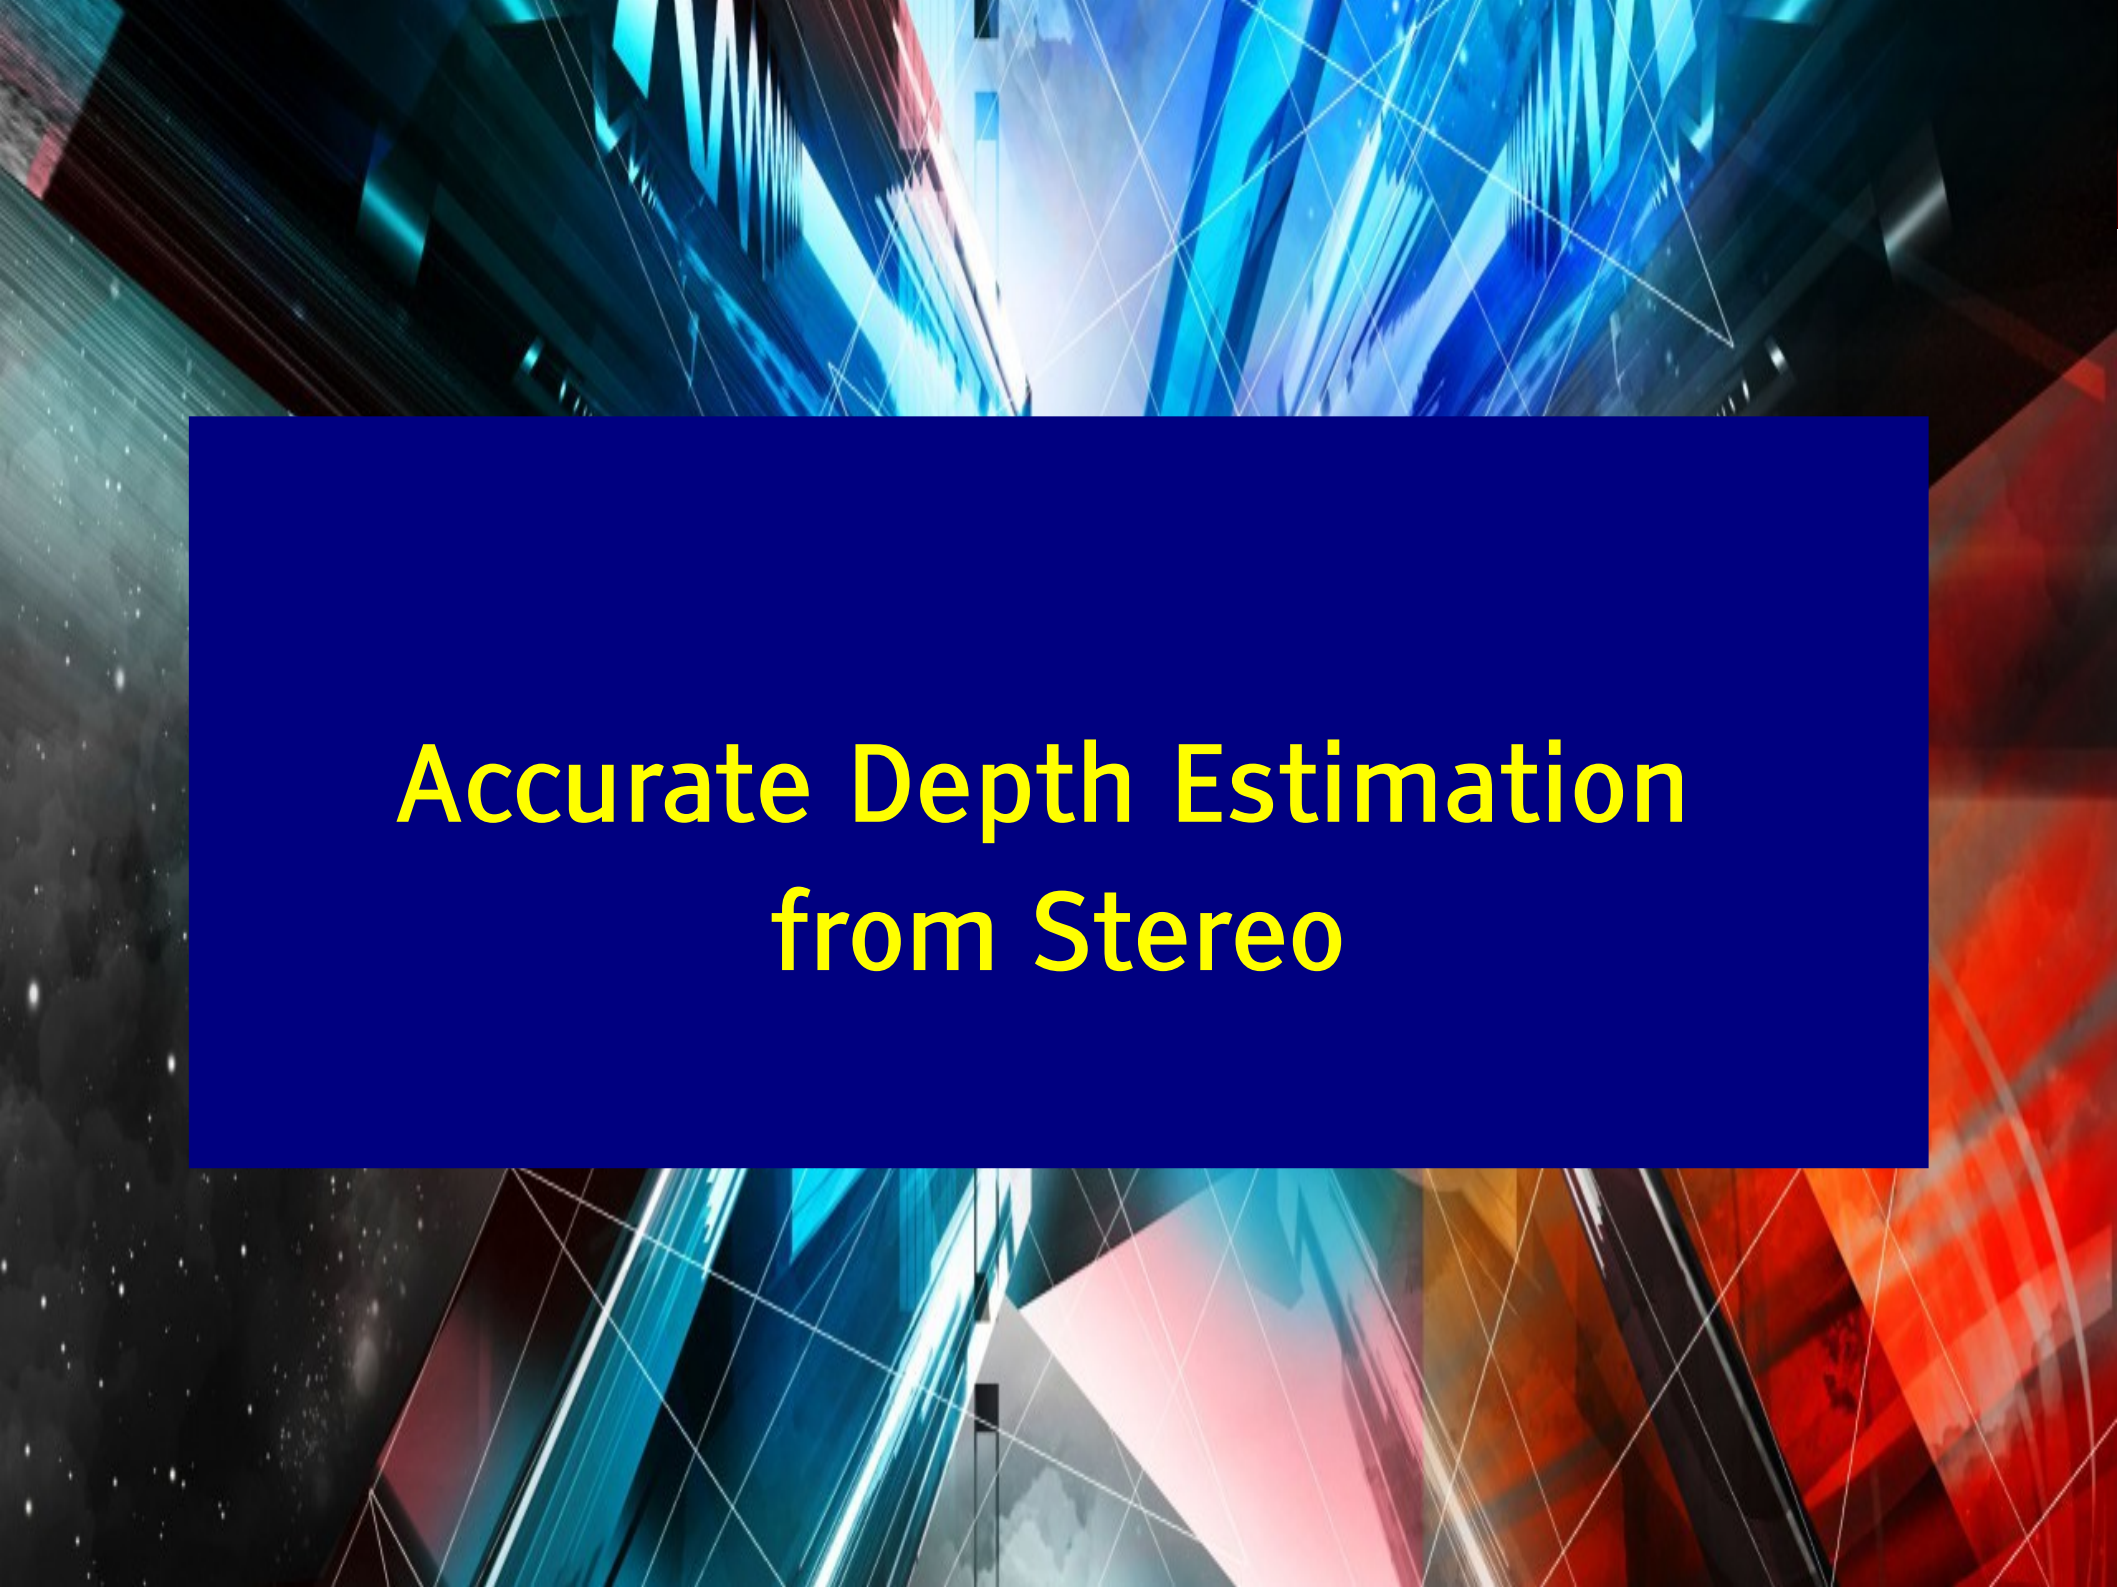
Accurate Depth Estimation
from Stereo

KITTI Dataset (RGB registered with LIDAR)

- Stereo Cameras + Velodyne LIDAR: aligned data.
- Collected from a car
- Lots of results with many methods from many teams around the world
- Stereo images sparsely labeled with depth values
- supervised learning

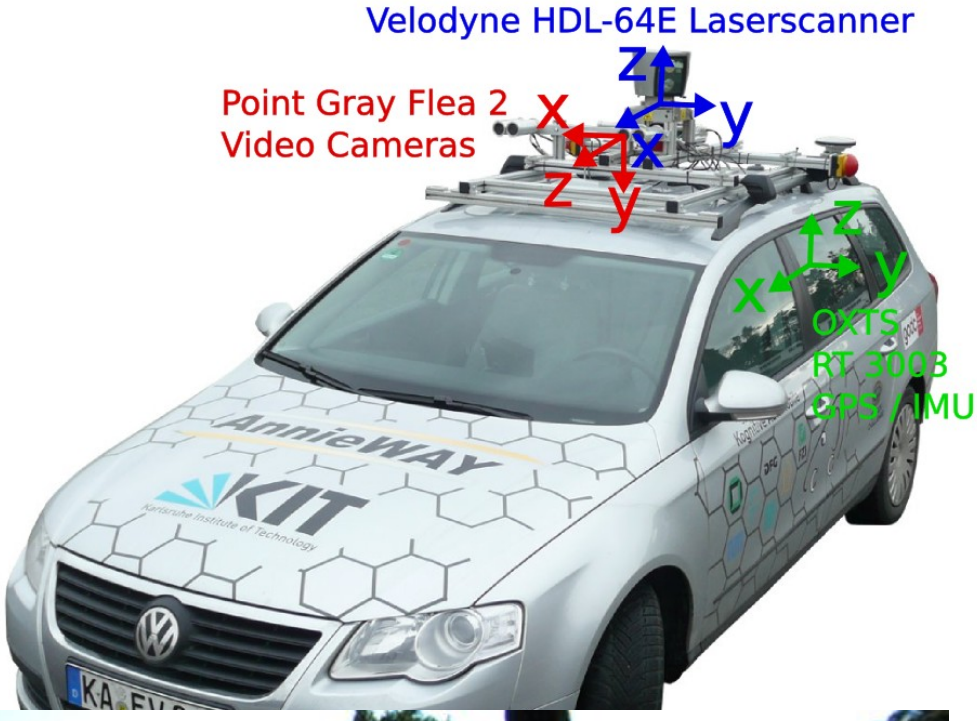

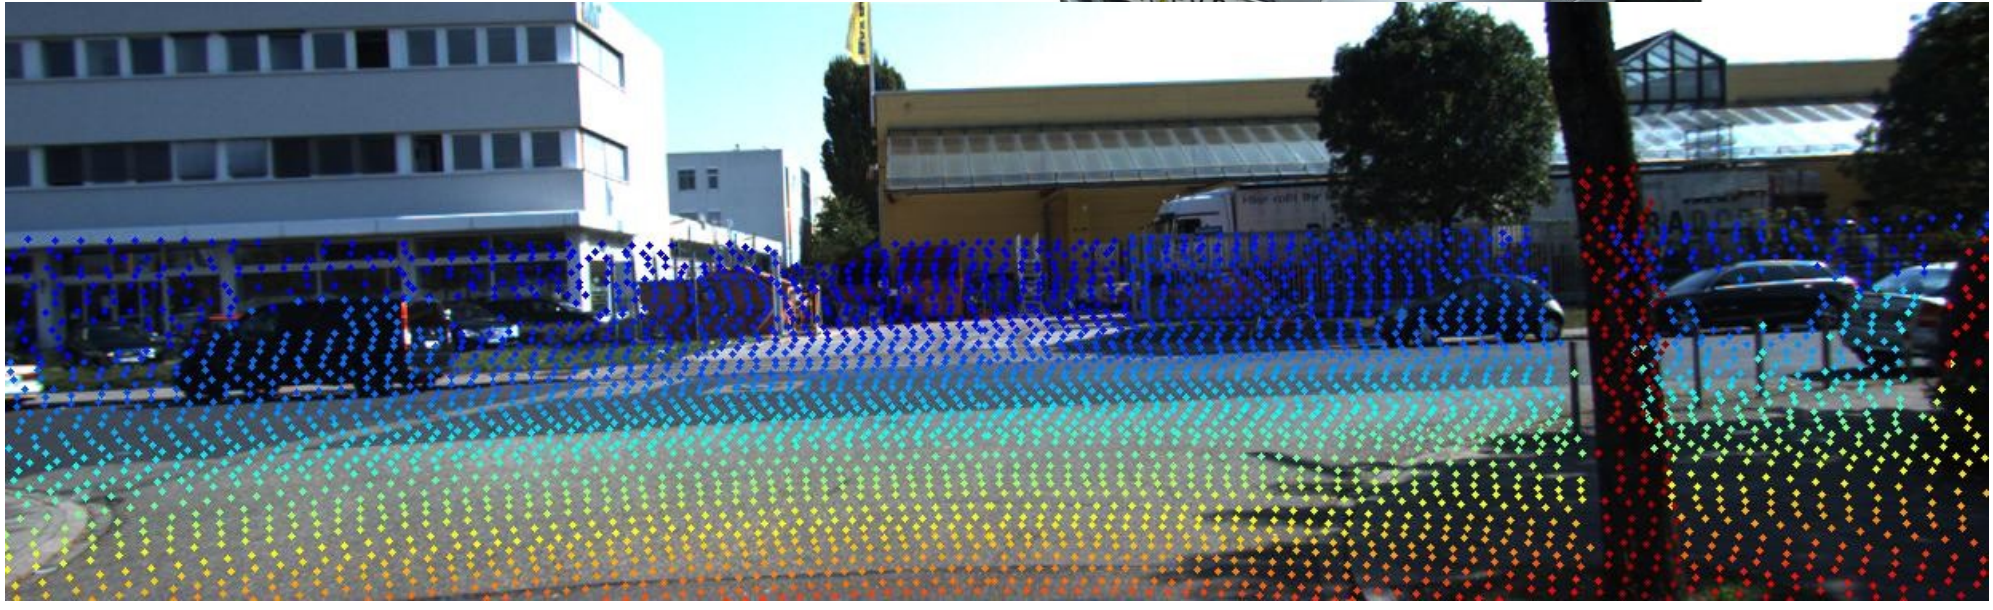

KITTI Dataset (RGB registered with LIDAR)

- supervised learning of a patch matcher (as a binary classifier)

Left patch	Right patch	Label
		Good match
		Bad match

ConvNet for Stereo Matching

Using a ConvNet to learn a similarity measure between image patches

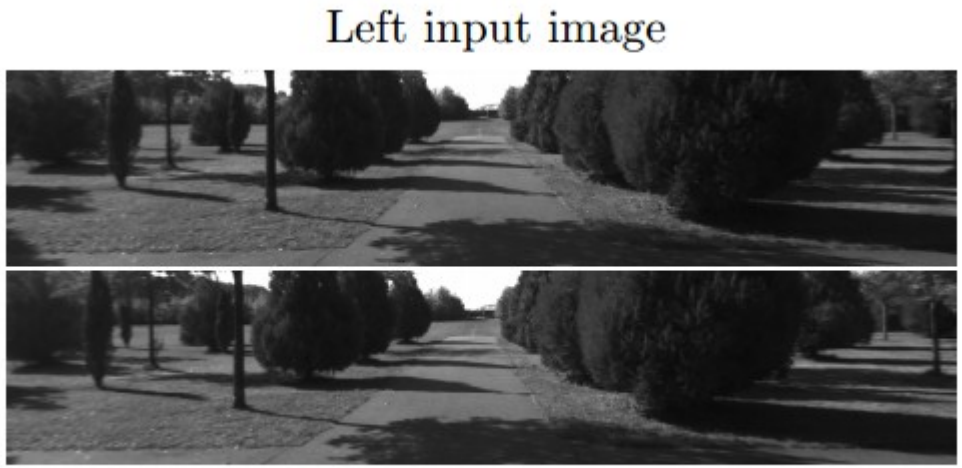

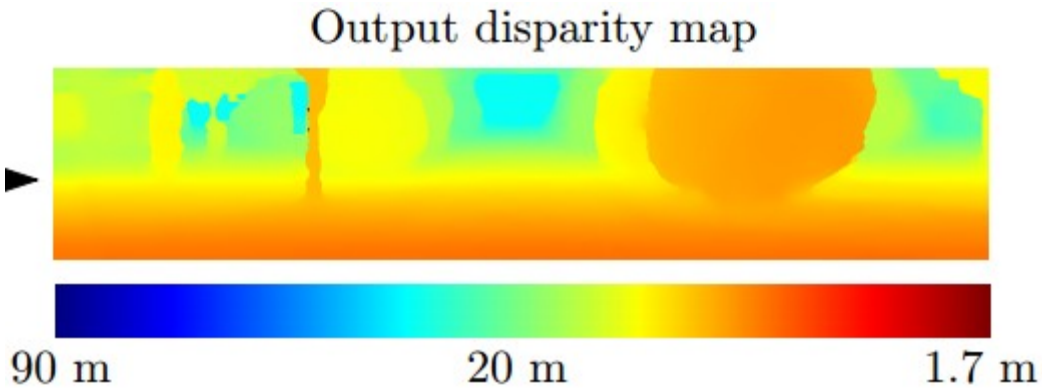

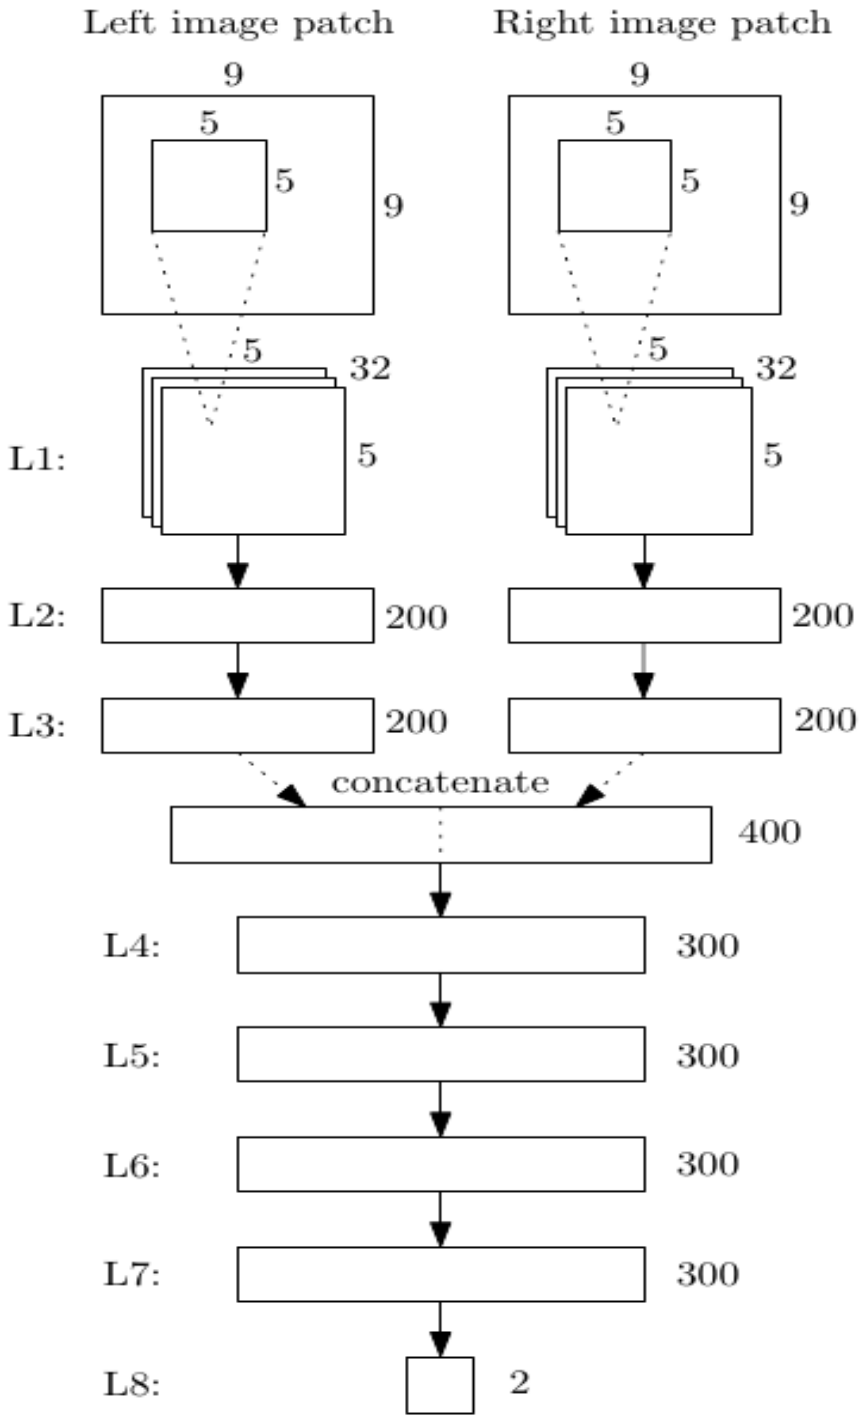

Figure 2: Network architecture

ConvNet for Stereo Matching

Image

Winner-take-all output
(14.55% error >3 pix)

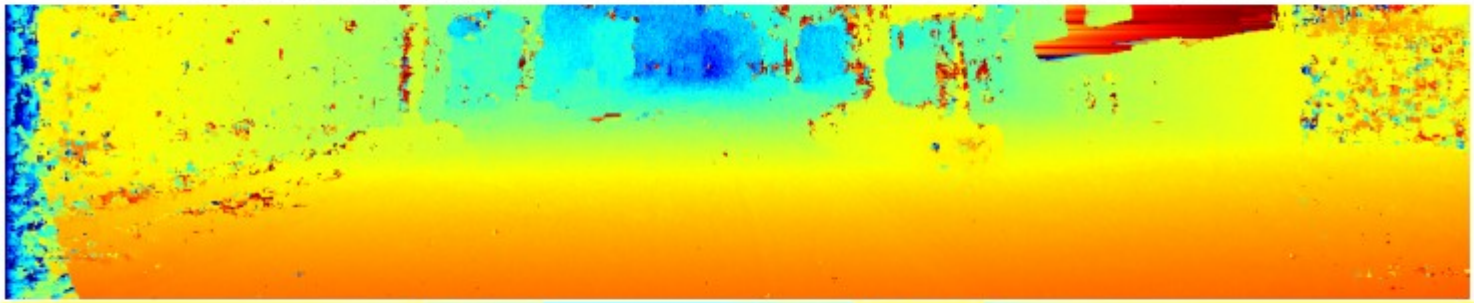

Result after cleanup
(2.61% error >3 pix)

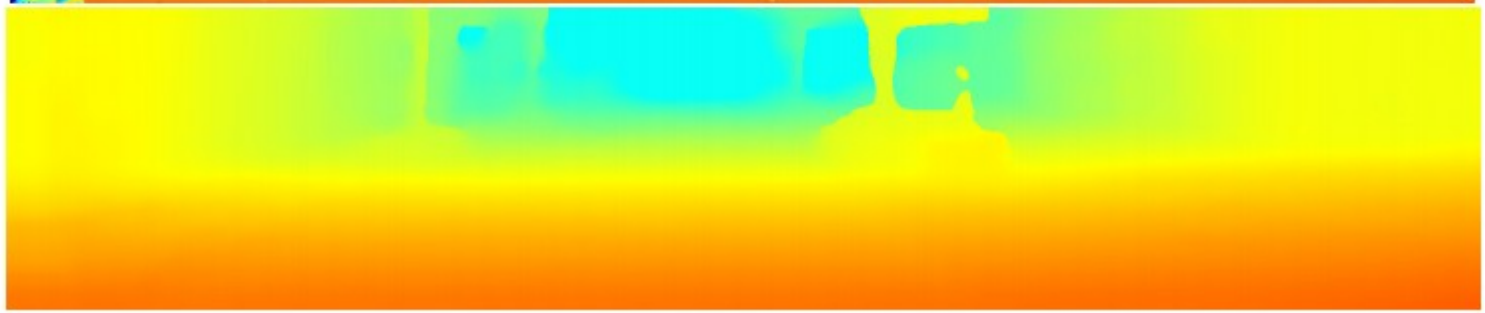

Depth Estimation from Stereo Pairs

Using a ConvNet to learn a similarity measure between image patches/

Record holder on KITTI dataset (Sept 2014):

Rank	Method		Error
1	Our method		2.61 %
2	SPS-StFl	Yamaguchi et al. (2014)	2.83 %
3	VC-SF	Vogel et al. (2014)	3.05 %
4	PCBP-SS	Yamaguchi et al. (2013)	3.40 %
5	DDS-SS		3.83 %

Original image

SADD

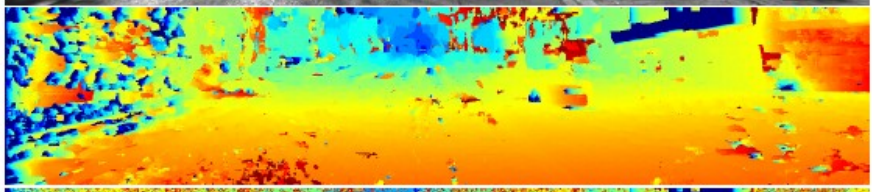

Census

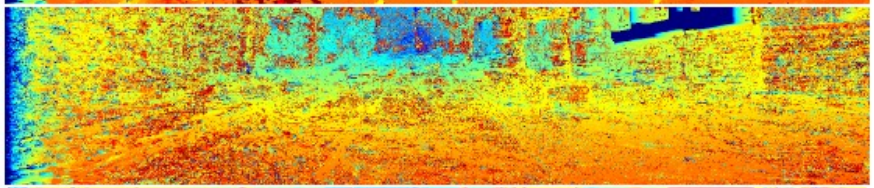

ConvNet

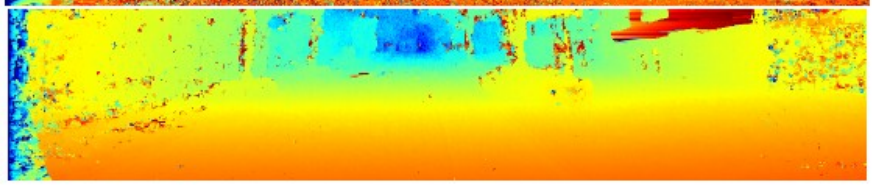

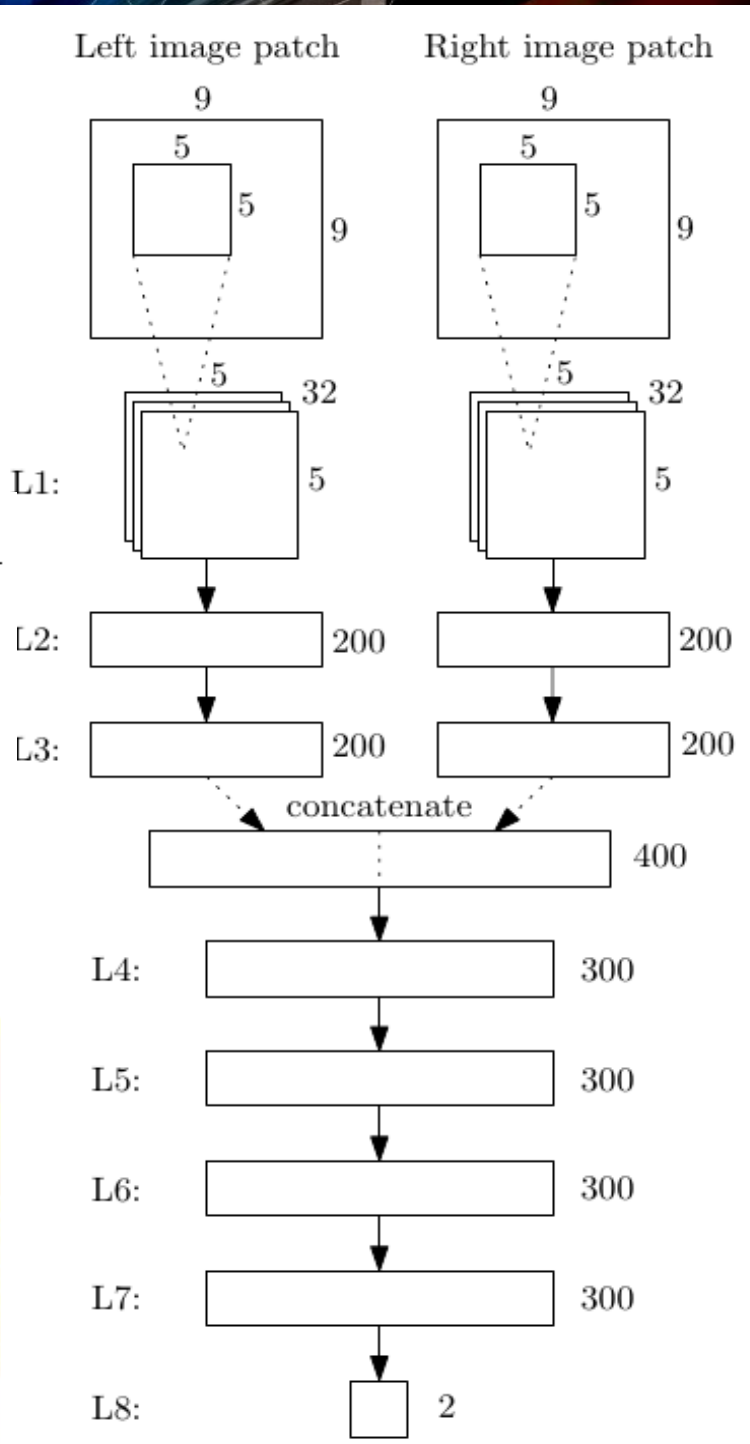

Figure 2: Network architecture

KITTI Stereo Leaderboard (Sept 2014)

Y LeCun

Metric: Percentage of pixels with more than 3 pixel error on disparity

Our ConvNet method is first. Doesn't use optical flow nor multiple images.

#2 uses optical flow and 2 frames

#3 uses multiple frames

#4 is at 3.39% vs our 2.61%

Rank	Method	Setting	Code	Out-Noc	Out-All	Avg-Noc	Avg-All	Density	Runtime
1	MC-CNN			2.61 %	3.84 %	0.8 px	1.0 px	100.00 %	100 s
J. Zbontar and Y. LeCun: Computing the Stereo Matching Cost with a Convolutional Neural Network. 2014.									
2	SPS-StFl	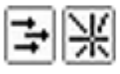		2.83 %	3.64 %	0.8 px	0.9 px	100.00 %	35 s
K. Yamaguchi, D. McAllester and R. Urtasun: Efficient Joint Segmentation, Occlusion Labeling, Stereo and Flow Estimation. ECCV 2									
3	VC-SF			3.05 %	**3.31 %**	**0.8 px**	0.8 px	100.00 %	300 s
C. Vogel, S. Roth and K. Schindler: View-Consistent 3D Scene Flow Estimation over Multiple Frames. Proceedings of European Conf									
4	SPS-St			3.39 %	4.41 %	0.9 px	1.0 px	100.00 %	5 s
K. Yamaguchi, D. McAllester and R. Urtasun: Efficient Joint Segmentation, Occlusion Labeling, Stereo and Flow Estimation. ECCV 2									
5	PCBP-SS			3.40 %	4.72 %	0.8 px	1.0 px	100.00 %	5 min
K. Yamaguchi, D. McAllester and R. Urtasun: Robust Monocular Epipolar Flow Estimation. CVPR 2013.									

Depth Estimation from Stereo Pairs: Results

[Zbontar & LeCun Arxiv '14]

Presentation at ECCV workshop 2014/9/6

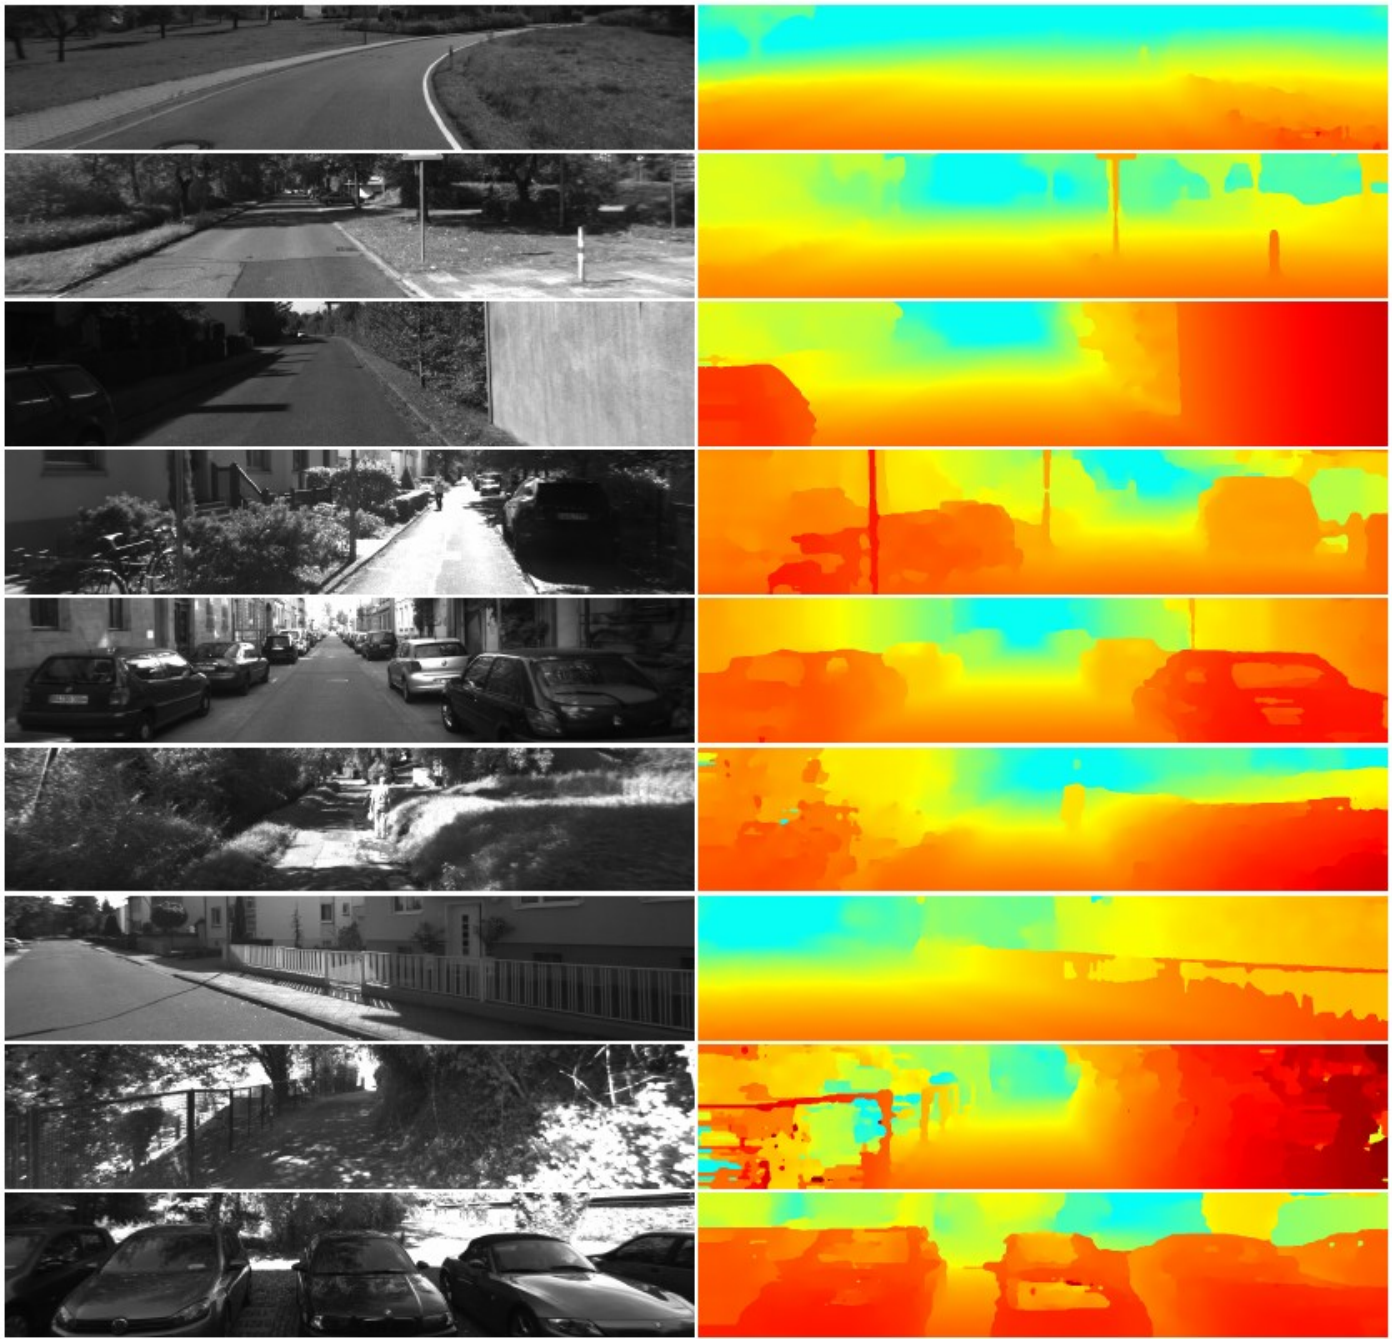

Body Pose Estimation

Pose Estimation and Attribute Recovery with ConvNets

Y LeCun

Pose-Aligned Network for Deep Attribute Modeling

[Zhang et al. CVPR 2014] (Facebook AI Research)

(a) Highest scoring results for people wearing glasses.

(b) Highest scoring results for people wearing a hat.

Real-time hand pose recovery

[Tompson et al. Trans. on Graphics 14]

Body pose estimation [Tompson et al. ICLR, 2014]

Other Tasks for Which Deep Convolutional Nets are the Best

- Handwriting recognition MNIST (many), Arabic HWX (IDSIA)
- OCR in the Wild [2011]: StreetView House Numbers (NYU and others)
- Traffic sign recognition [2011] GTSRB competition (IDSIA, NYU)
- Asian handwriting recognition [2013] ICDAR competition (IDSIA)
- Pedestrian Detection [2013]: INRIA datasets and others (NYU)
- Volumetric brain image segmentation [2009] connectomics (IDSIA, MIT)
- Human Action Recognition [2011] Hollywood II dataset (Stanford)
- Object Recognition [2012] ImageNet competition (Toronto)
- Scene Parsing [2012] Stanford bgd, SiftFlow, Barcelona datasets (NYU)
- Scene parsing from depth images [2013] NYU RGB-D dataset (NYU)
- Speech Recognition [2012] Acoustic modeling (IBM and Google)
- Breast cancer cell mitosis detection [2011] MITOS (IDSIA)

- The list of perceptual tasks for which ConvNets hold the record is growing.
- Most of these tasks (but not all) use purely supervised convnets.

A Few Vision Projects
At FAIR

Video Classification

Video Classification using 3D Convolutional Networks:

- CVPR 2015 submission with Du Tran (Dartmouth) , Lubomir, Rob, Manohar
- Train 3D ConvNets on Instagram Videos
- Use as feature extractor (C3D) for other video tasks/datasets

- State-of-the-art results across several benchmarks for object, scene, action tasks.
- E.g. beats Google & Stanford on Sports 1M (44% AP vs 41.5%, & still training…)

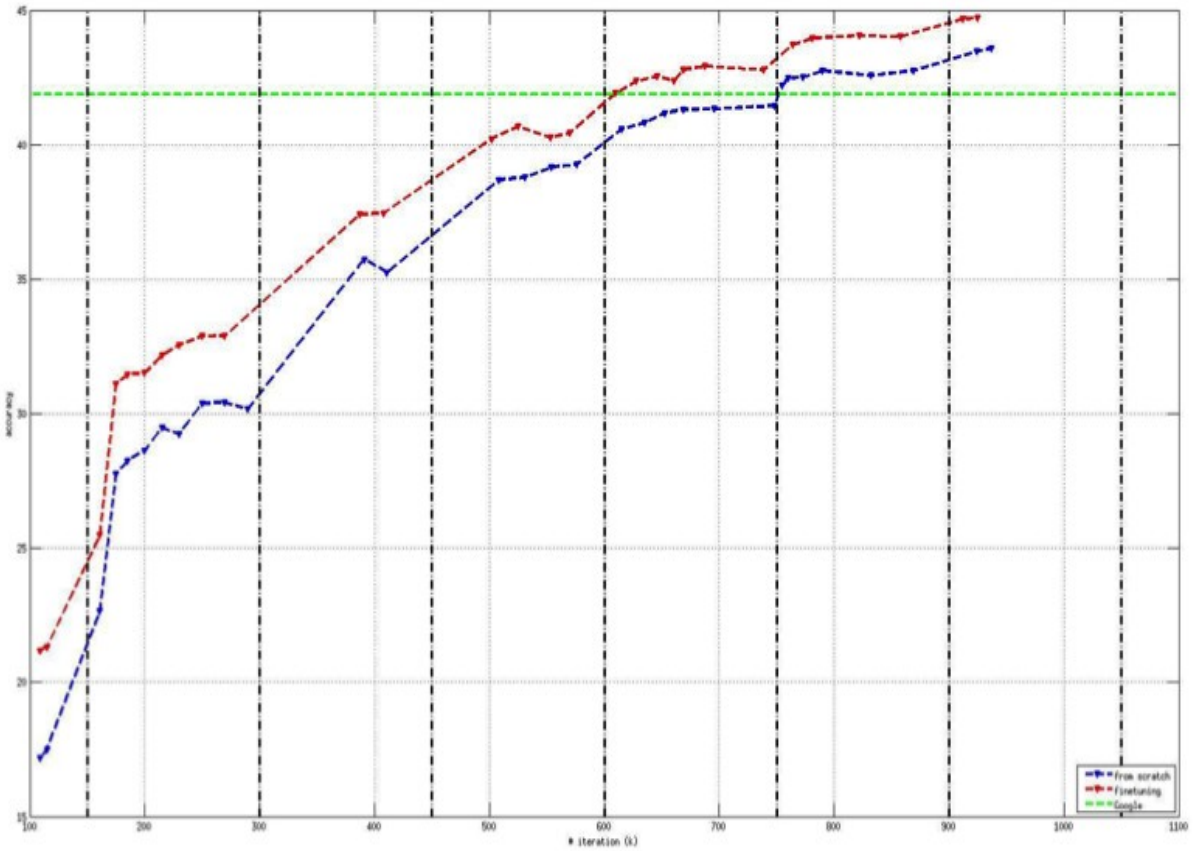

Compact Binary Codes

Work by Yunchao Gong

- Many applications throughout FB

1. Detecting copyrighted videos

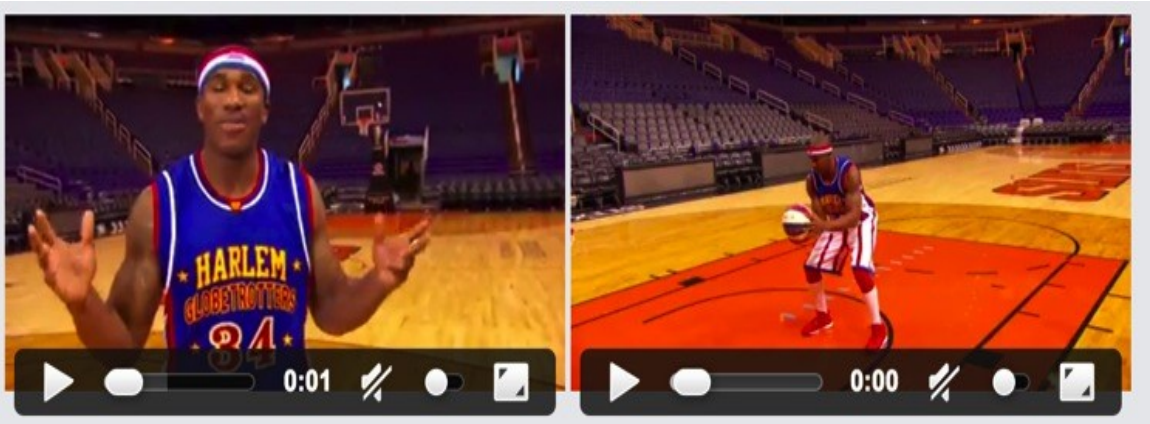

Same video or not?

2. Large-scale image search

3. Finding memes for spam detection:

4. online semantic clustering for event detection:

Unsupervised Learning from Video

- ICLR 2015 submission with Marc'Aurelio Ranzato, Arthur Szlam, Joan Bruna, Michael Matheiu, Ronan Collobert, Sumit Chopra

- Use language models (RNN etc.)
  to model video patches

- Learn good representations without need for labels

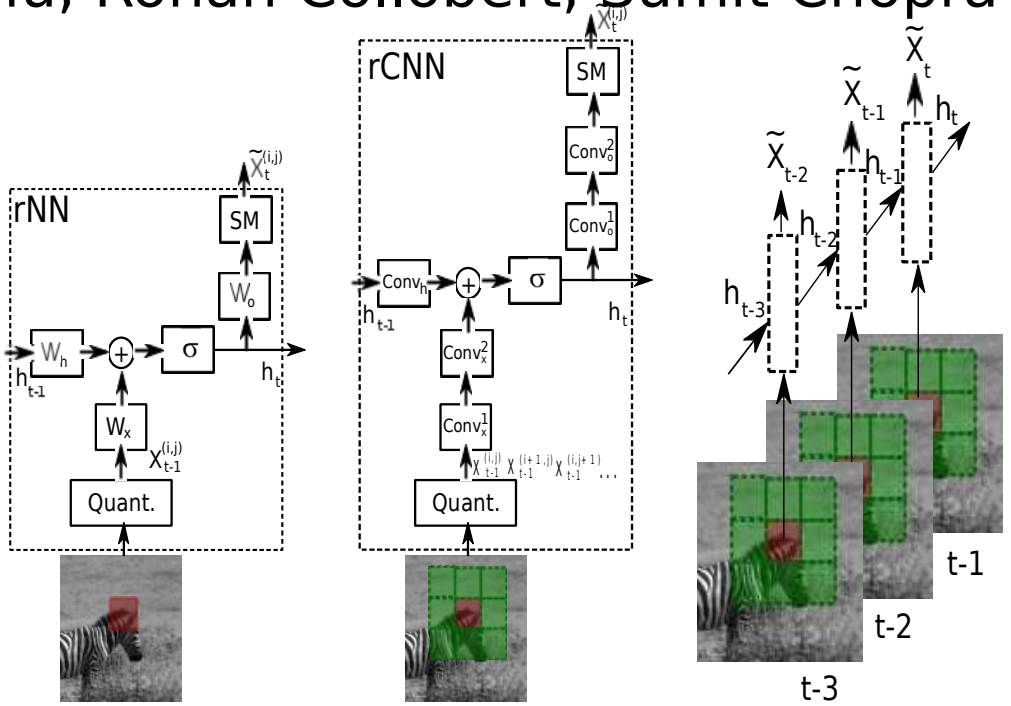

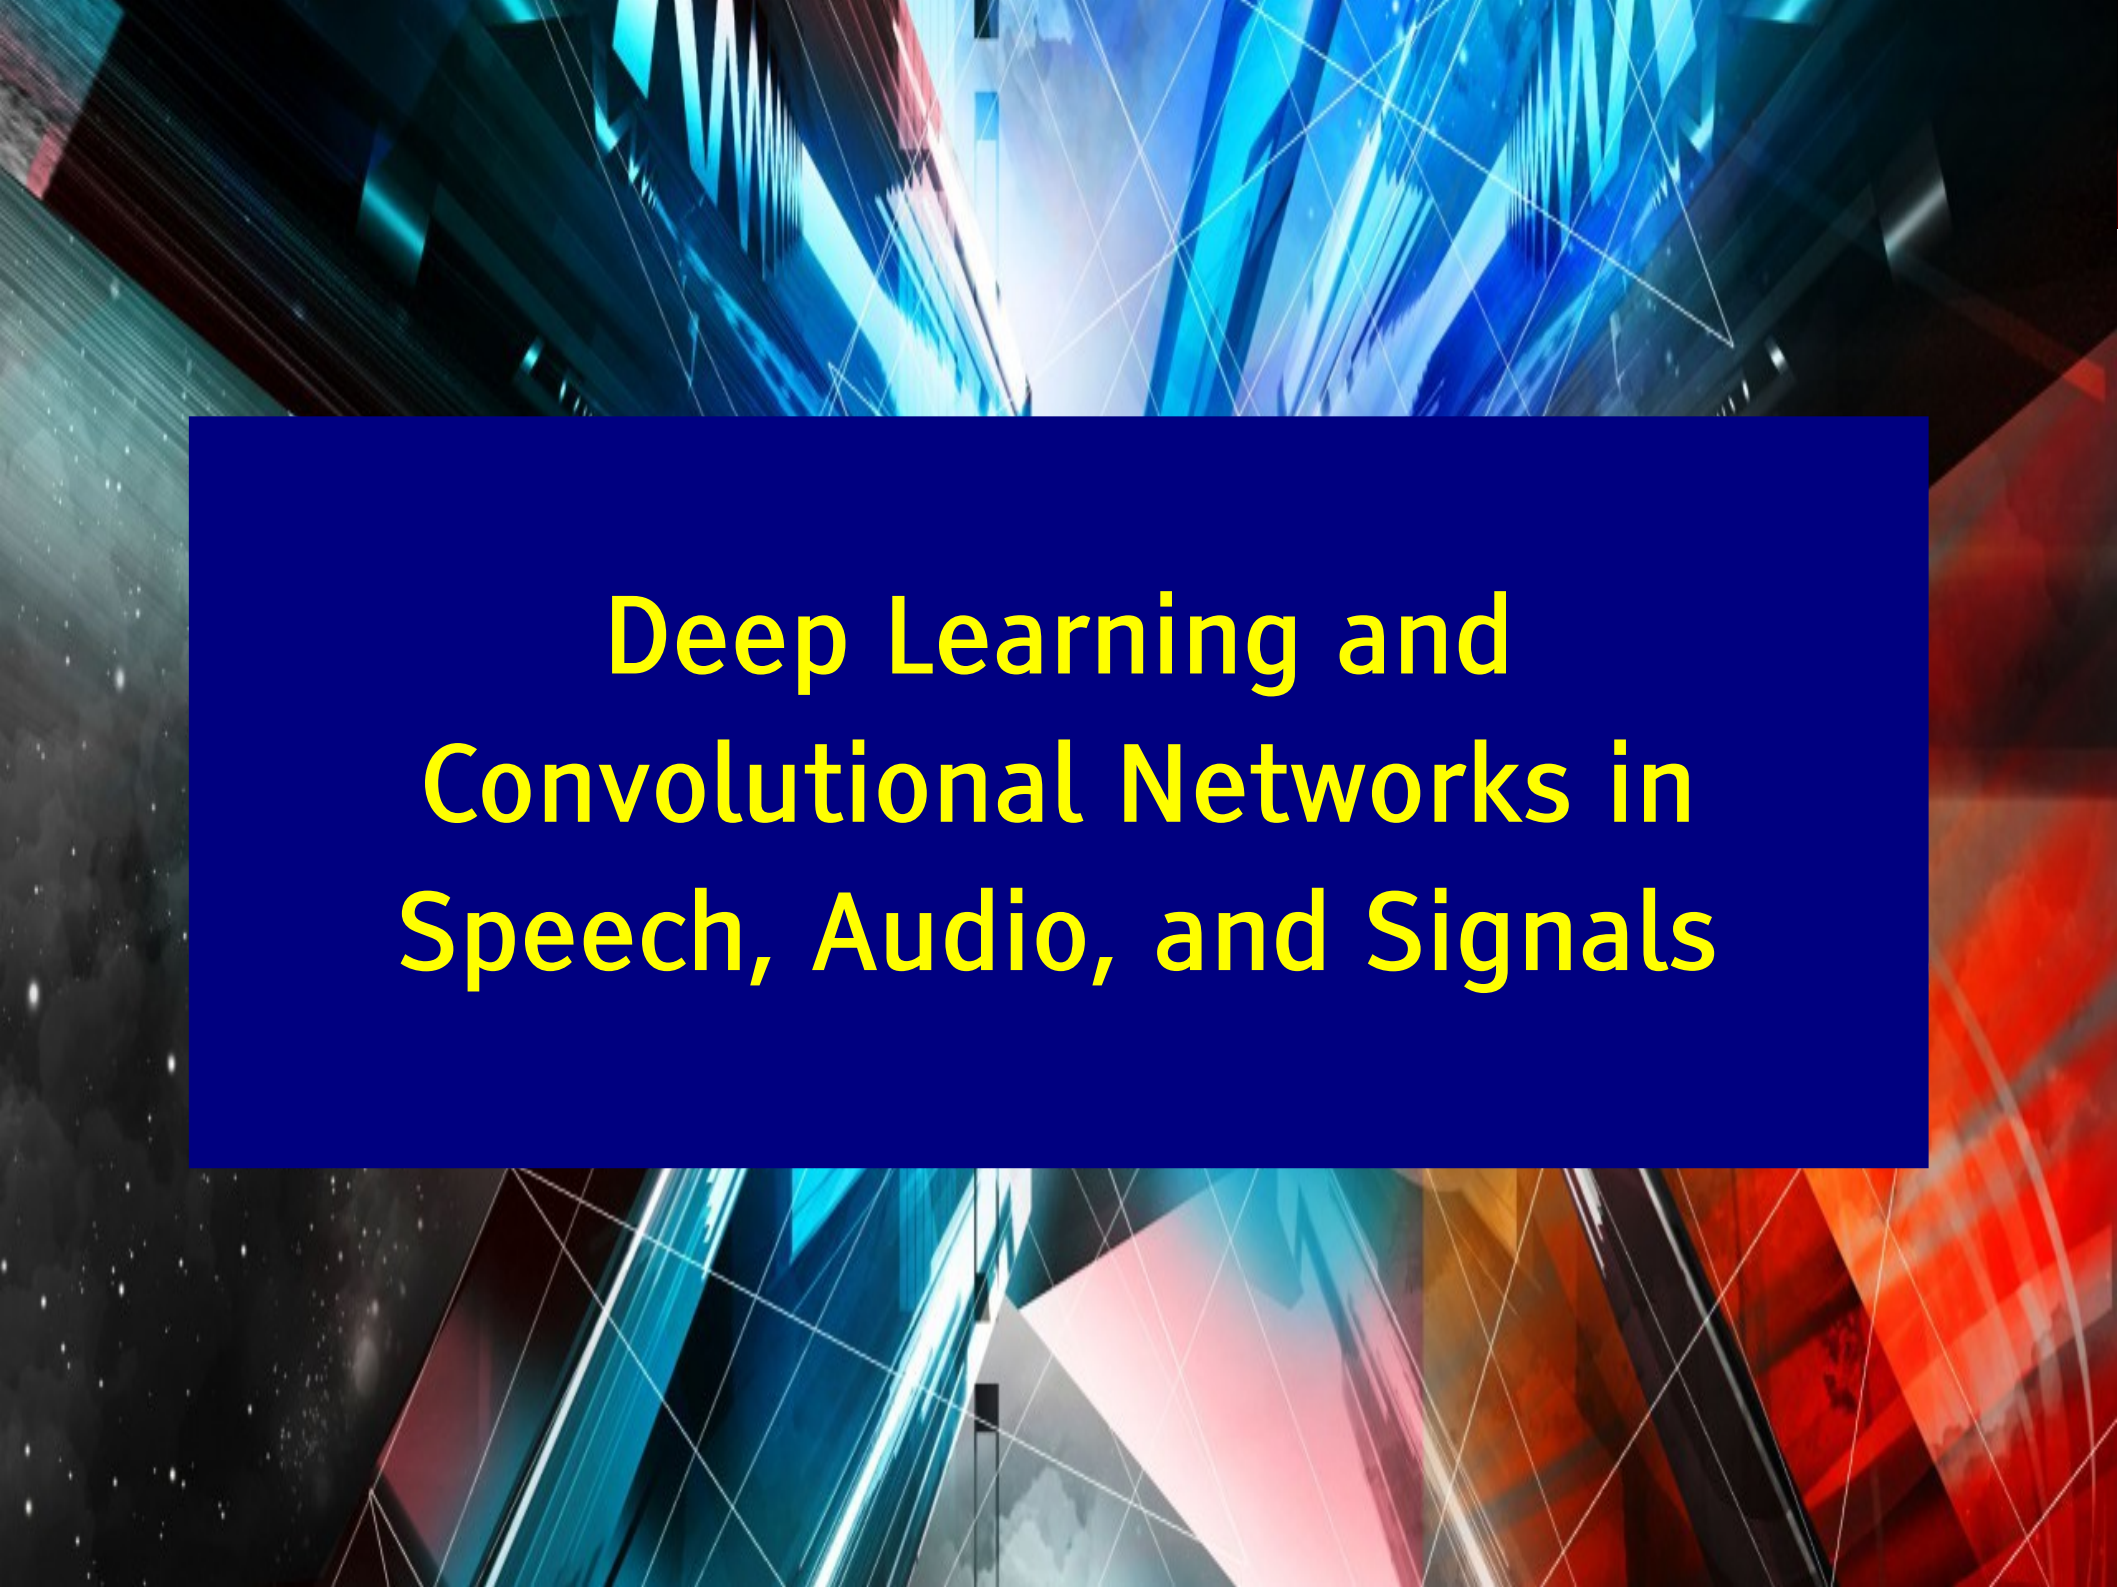
Deep Learning and
Convolutional Networks in
Speech, Audio, and Signals

Acoustic Modeling in Speech Recognition (Google)

- A typical speech recognition architecture with DL-based acoustic modeling
  - Features: log energy of a filter bank (e.g. 40 filters)
  - Neural net acoustic modeling (convolutional or not)
  - Input window: typically 10 to 40 acoustic frames
  - Fully-connected neural net: 10 layers, 2000-4000 hidden units/layer
  - But convolutional nets do better....
  - Predicts phone state, typically 2000 to 8000 categories

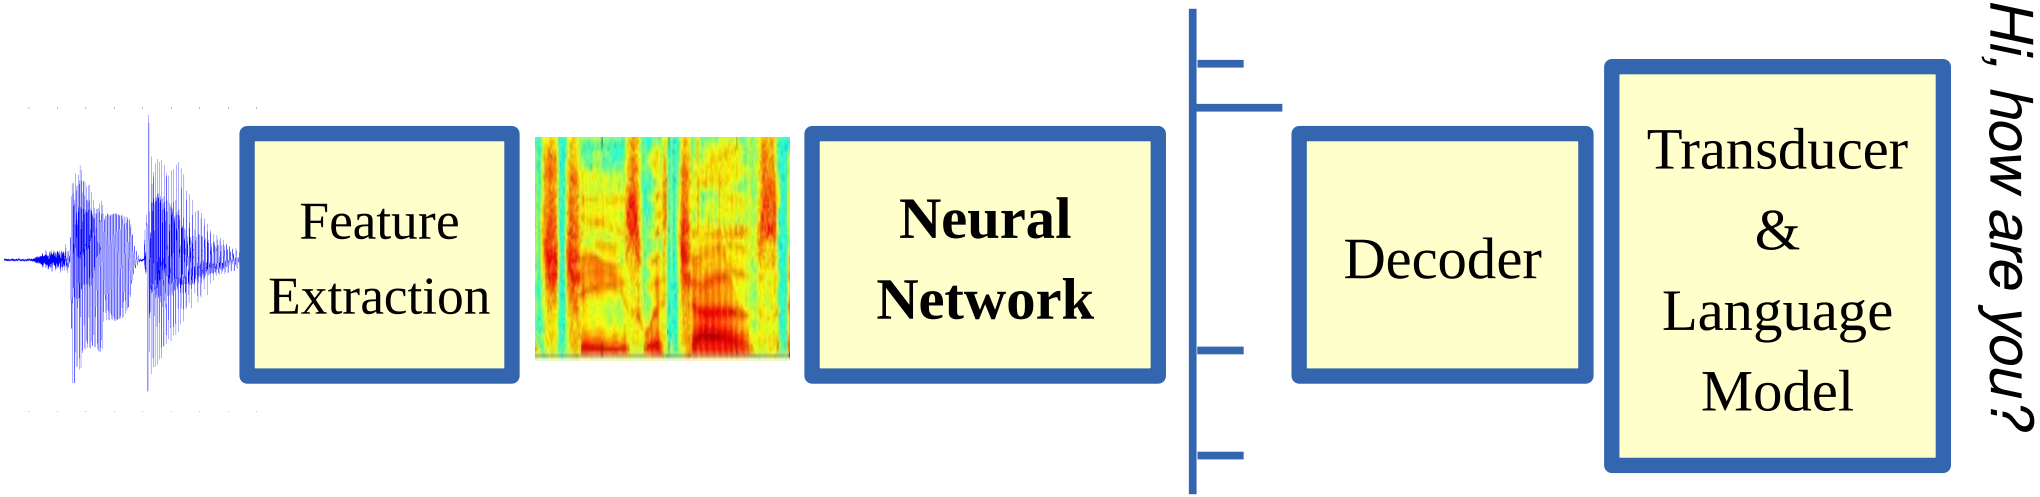

Mohamed et al. “DBNs for phone recognition” NIPS Workshop 2009
Zeiler et al. “On rectified linear units for speech recognition” ICASSP 2013

Speech Recognition with Convolutional Nets (NYU/IBM)

Y LeCun

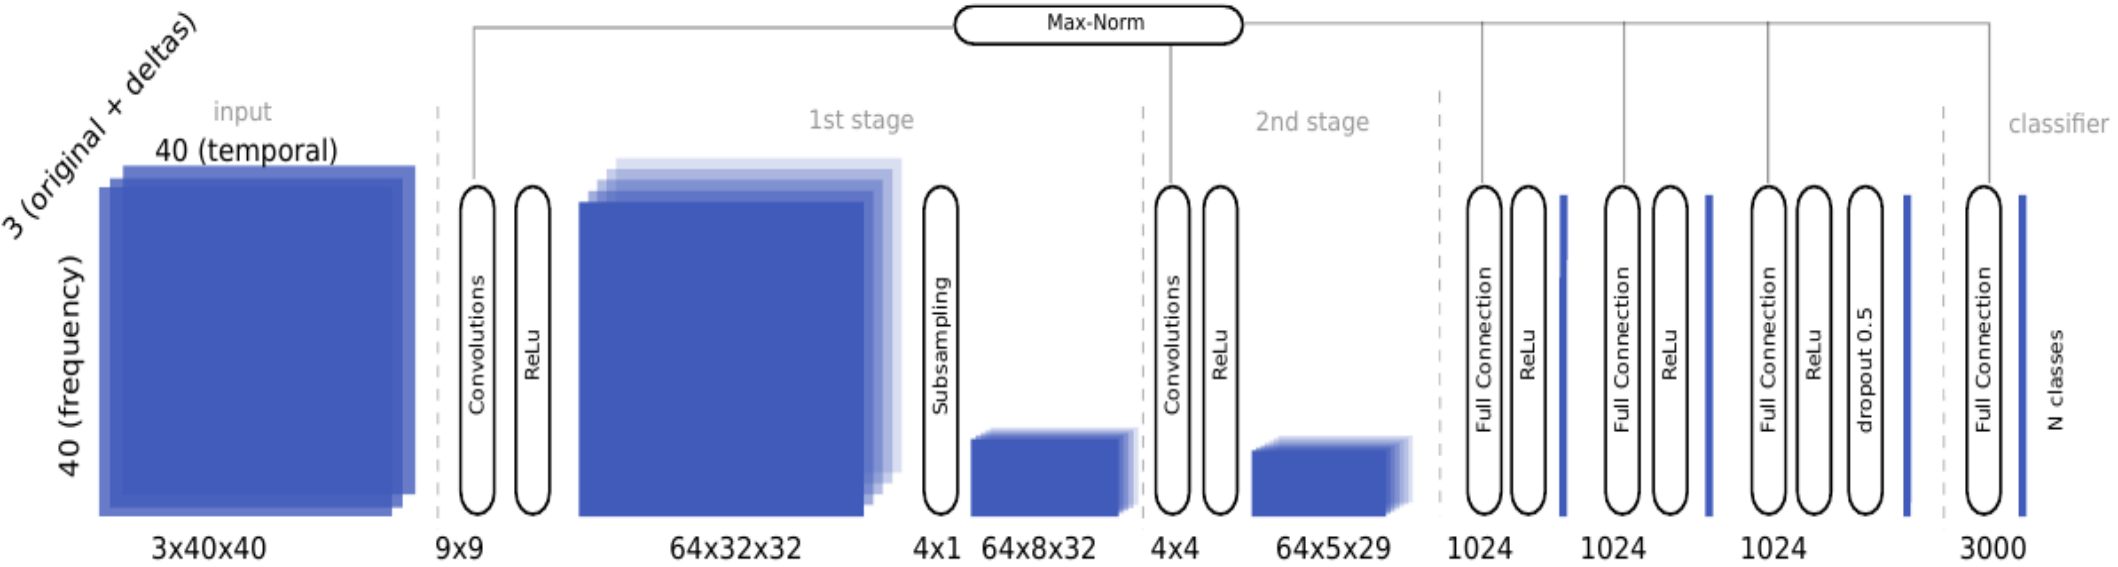

- Acoustic Model: ConvNet with 7 layers. 54.4 million parameters.
- Classifies acoustic signal into 3000 context-dependent subphones categories
- ReLU units + dropout for last layers
- Trained on GPU. 4 days of training

Speech Recognition with Convolutional Nets (NYU/IBM)

Y LeCun

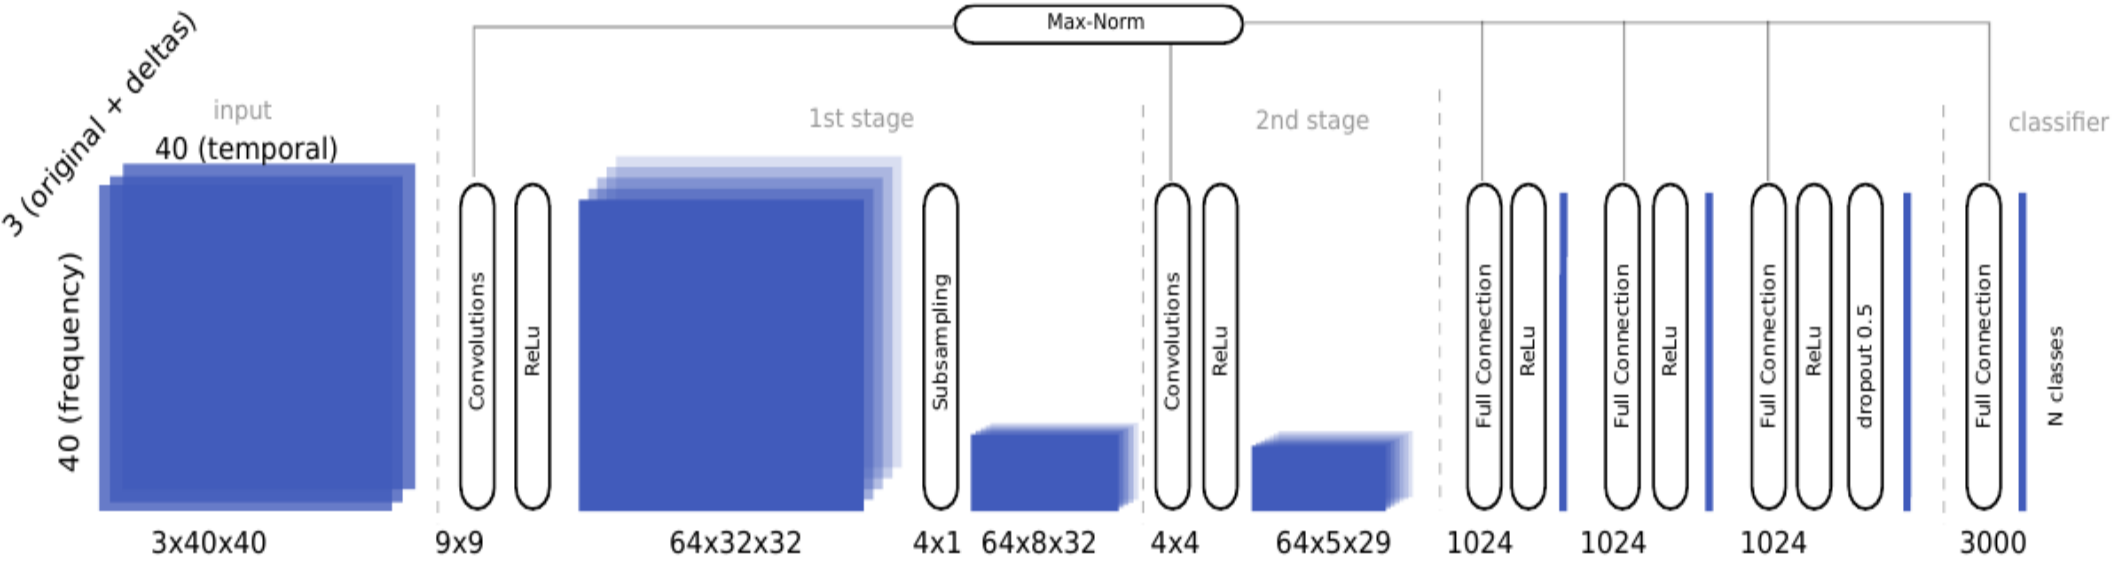

- **Subphone-level classification error (sept 2013):**
  - Cantonese: phone: 20.4% error; subphone: 33.6% error (IBM DNN: 37.8%)
- **Subphone-level classification error (march 2013)**
  - Cantonese: subphone: 36.91%
  - Vietnamese: subphone 48.54%
  - Full system performance (token error rate on conversational speech):
    - 76.2% (52.9% substitution, 13.0% deletion, 10.2% insertion)

Speech Recognition with Convolutional Nets (NYU/IBM)

- Training samples.
  - 40 MEL-frequency Cepstral Coefficients
  - Window: 40 frames, 10ms each

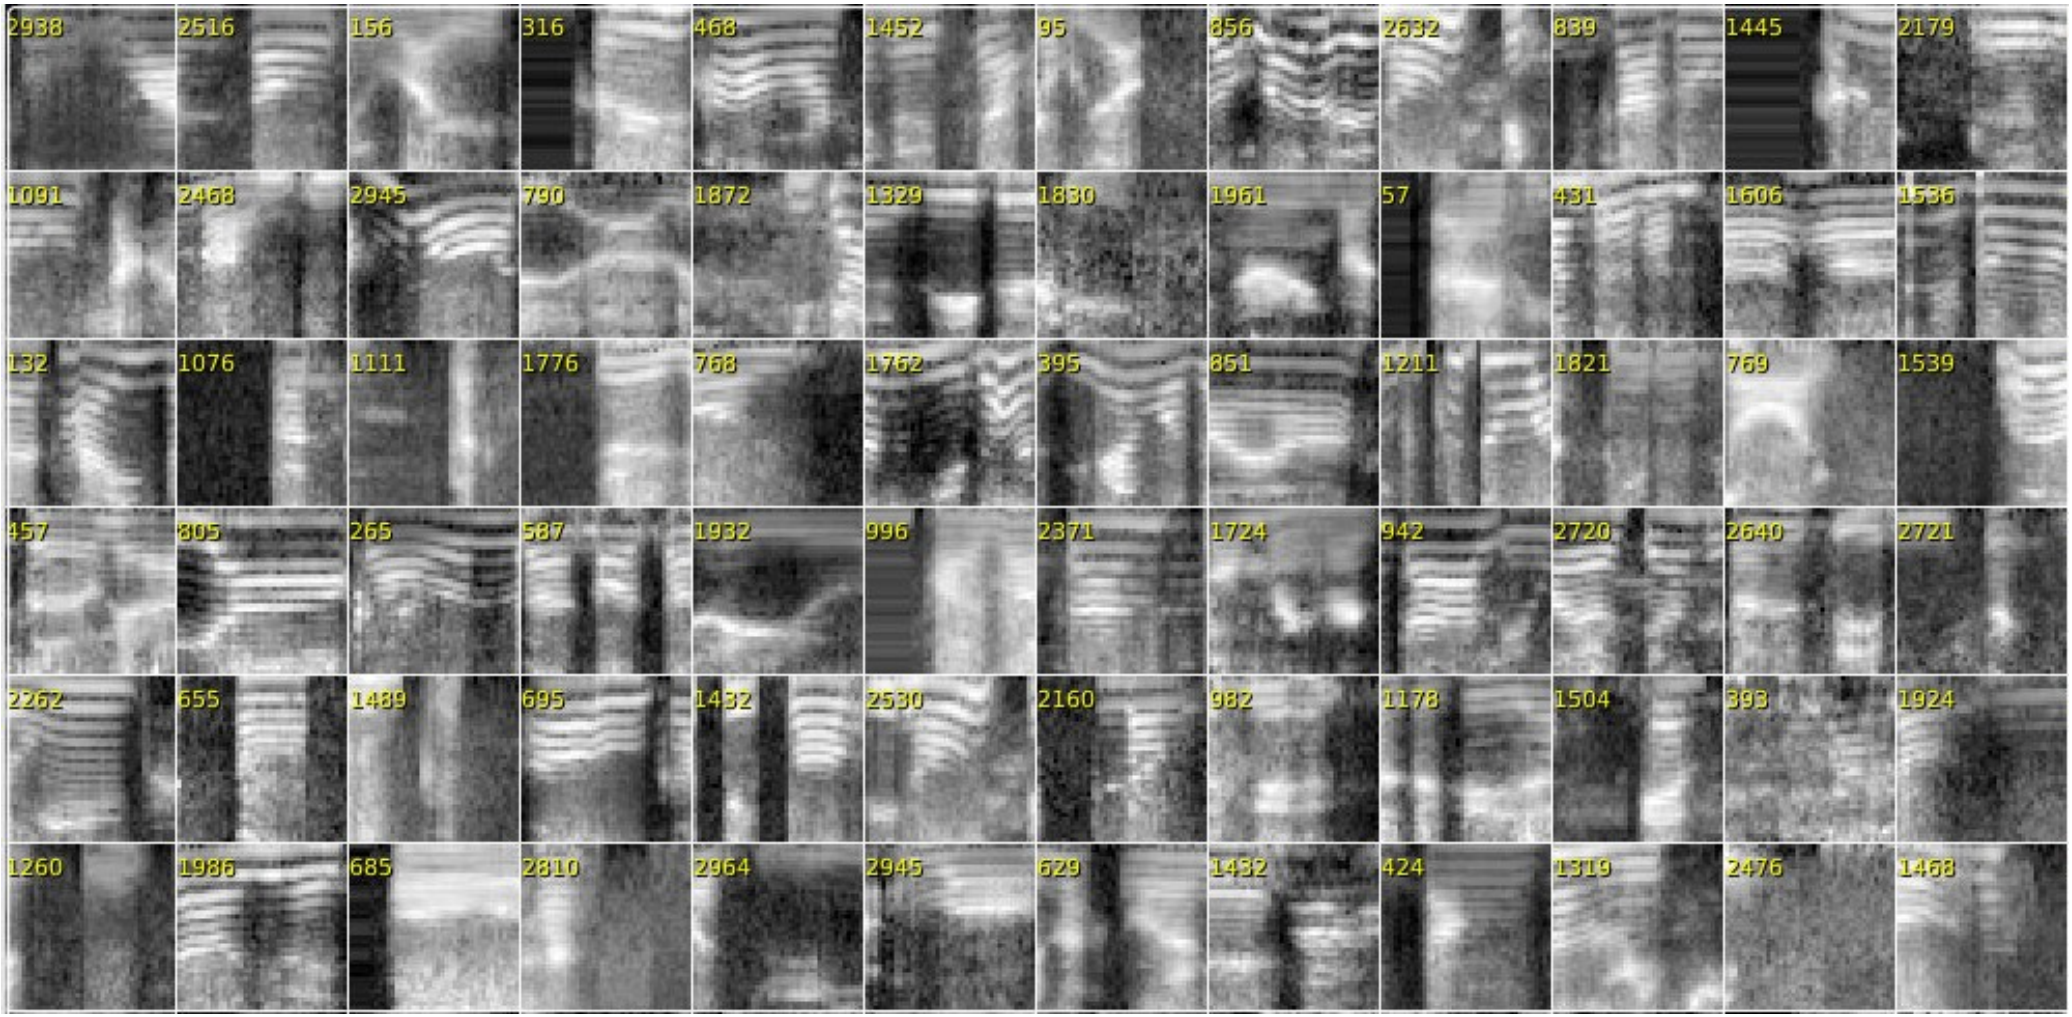

Speech Recognition with Convolutional Nets (NYU/IBM)

- Convolution Kernels at Layer 1:
  - 64 kernels of size 9x9

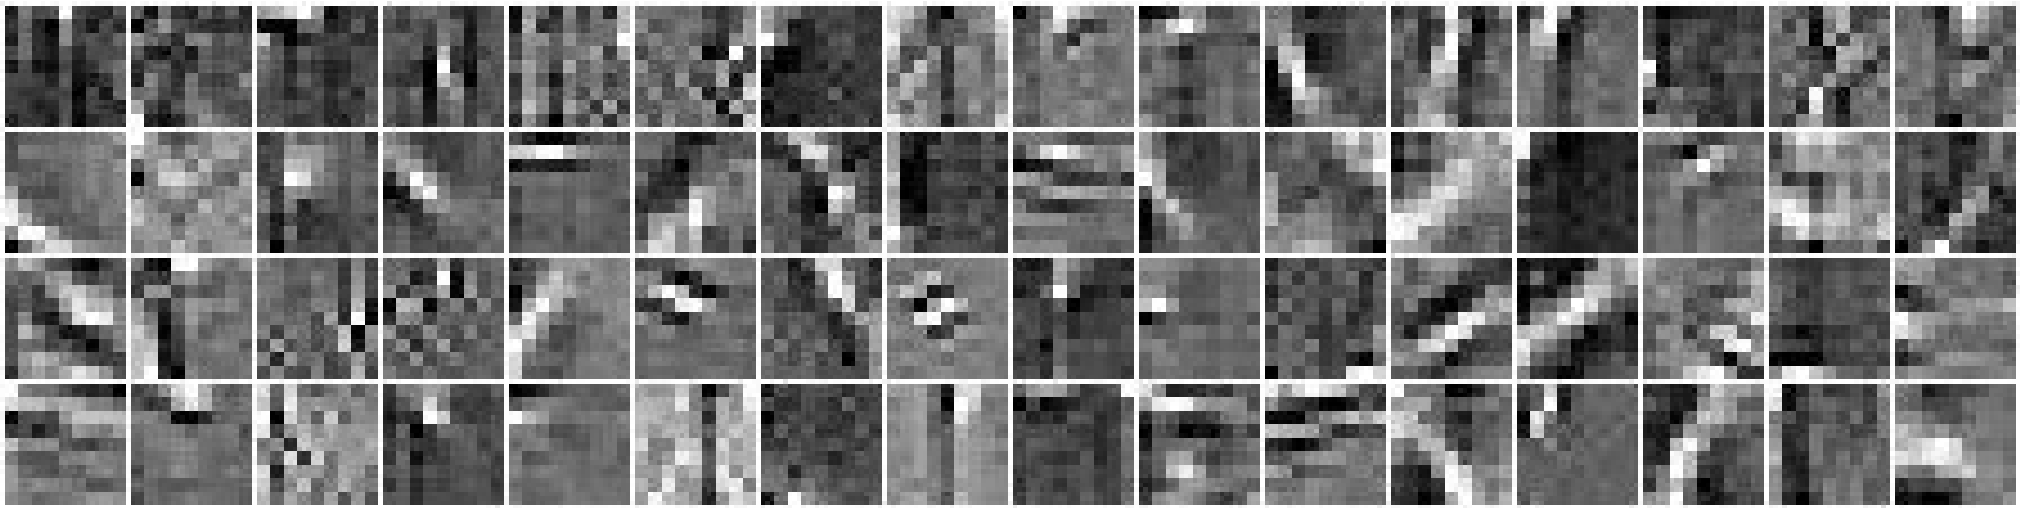

Software Tools
and
Hardware Acceleration
for
Convolutional Networks

Software Platform for Deep Learning: Torch7

- **Torch7**
  - based on the **LuaJIT** language
  - Simple and lightweight dynamic language (widely used for games)
  - Multidimensional array library with CUDA and OpenMP backends
  - FAST: Has a native just-in-time compiler
  - Has an unbelievably nice foreign function interface to call C/C++ functions from Lua
- **Torch7 is an extension of Lua with**
  - Multidimensional array engine
  - A machine learning library that implements multilayer nets, convolutional nets, unsupervised pre-training, etc
  - Various libraries for data/image manipulation and computer vision
  - Used at Facebook Ai Research, Google (Deep Mind, Brain), Intel, and many academic groups and startups
- **Single-line installation on Ubuntu and Mac OSX:**
  - http://torch.ch
- **Torch7 Cheat sheet (with links to libraries and tutorials):**
  - https://github.com/torch/torch7/wiki/Cheatsheet

Example: building a Neural Net in Torch7

- Net for SVHN digit recognition
- 10 categories
- Input is 32x32 RGB (3 channels)
- 1500 hidden units

```
Noutputs = 10;
nfeats = 3; Width = 32; height = 32
ninputs = nfeats*width*height
nhiddens = 1500
```

- Creating a 2-layer net
- Make a cascade module
- Reshape input to vector
- Add Linear module
- Add tanh module
- Add Linear Module
- Add log softmax layer

```
-- Simple 2-layer neural network
model = nn.Sequential()
model:add(nn.Reshape(ninputs))
model:add(nn.Linear(ninputs,nhiddens))
model:add(nn.Tanh())
model:add(nn.Linear(nhiddens,noutputs))
model:add(nn.LogSoftMax())
```

- Create loss function module

```
criterion = nn.ClassNLLCriterion()
```

See Torch7 example at http://bit.ly/16tyLAx

Recognizing Compositional Objects,
Structured Prediction,
Energy-Based Models

End-to-End Learning – Word-Level Discriminative Training

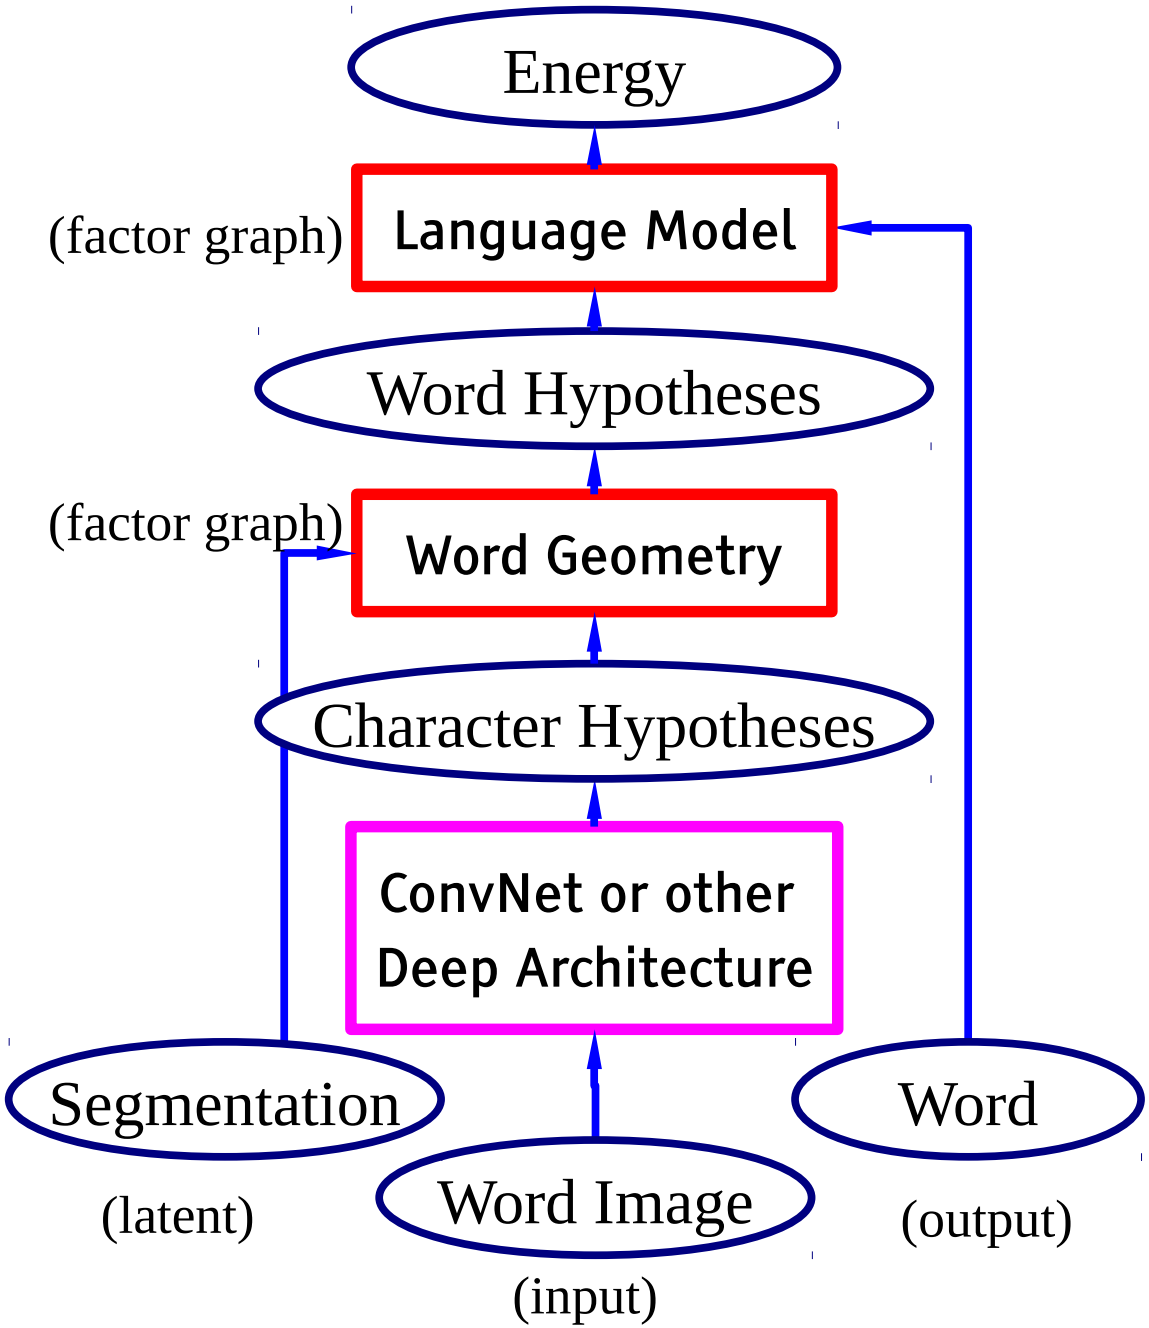

Making every single module in the system trainable.

Every module is trained simultaneously so as to optimize a global loss function.

Includes the feature extractor, the recognizer, and the contextual post-processor (graphical model)

Problem: back-propagating gradients through the graphical model.

Integrating Deep Learning and Structured Prediction

- Deep Learning systems can be assembled into factor graphs (graphical models)
  - Energy function measures the "compatibility" between variables.
  - X: inputs (observed variables)
  - Z: latent variables (never observed)
  - Y: outputs (observed on training set)
- Inference is energy minimization (MAP) or free energy minimization (marginalization) over Z and Y given an X
- Y* = Argmin E(X,Y,Z)
- Energy function can embed whole deep learning systems (e.g. ConvNets).

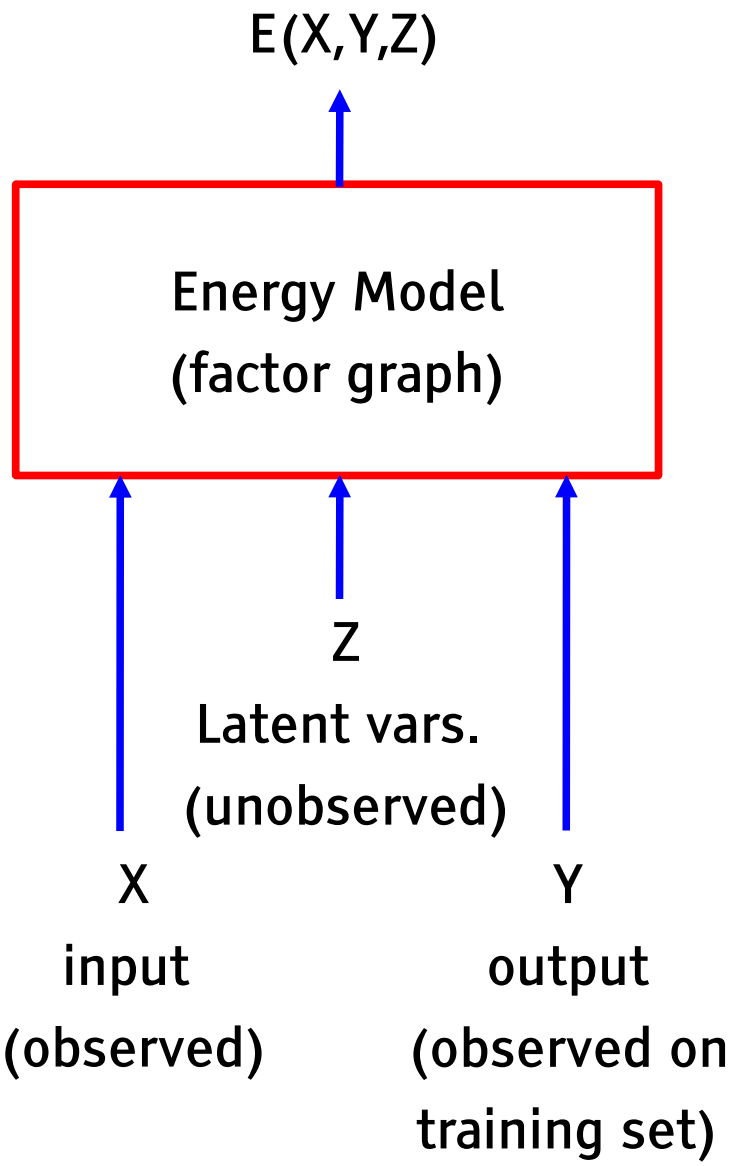

Latent Variable Models

The energy includes "hidden" variables Z whose value is never given to us

- We can minimize the energy over those latent variables
- We can also "marginalize" the energy over the latent variables

Minimization over latent variables:

$$E(Y,X) = \min_{Z \in \mathcal{Z}} E(Z,Y,X).$$

Marginalization over latent variables:

$$E(X,Y) = -\frac{1}{\beta} \log \int_{z \in \mathcal{Z}} e^{-\beta E(z,Y,X)}$$

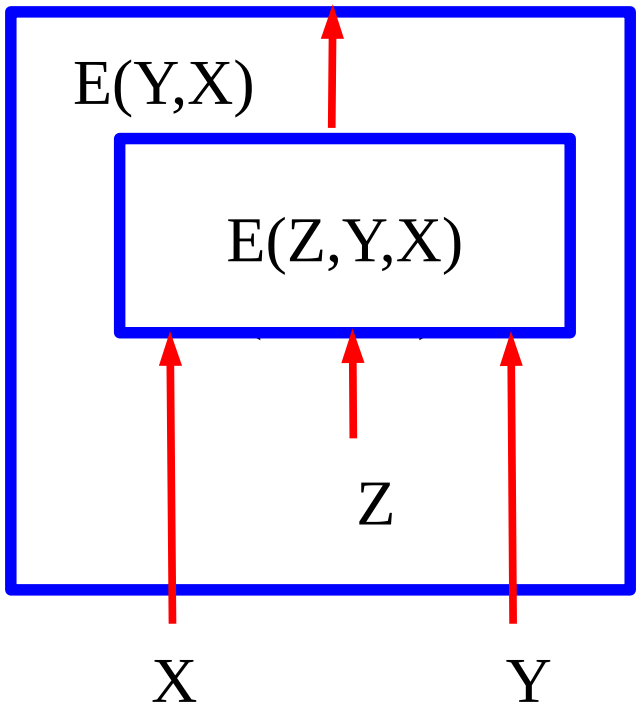

Estimating this integral may require some approximations (sampling, variational methods, maximum a posteriori....)

Loss Functions for EBM

Loss Functions

- Energy Loss (pull down on the data points)

$$L_{energy}(Y^i, E(W, \mathcal{Y}, X^i)) = E(W, Y^i, X^i).$$

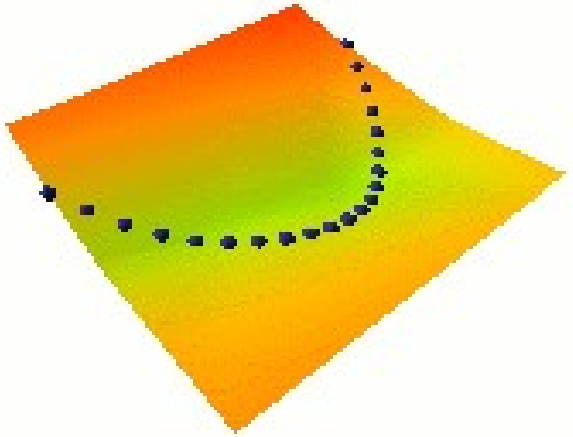

- Square-Square Loss

$$L_{\text{sq-sq}}(W, Y^i, X^i) = E(W, Y^i, X^i)^2 + \left(\max(0, m - E(W, \bar{Y}^i, X^i))\right)^2.$$

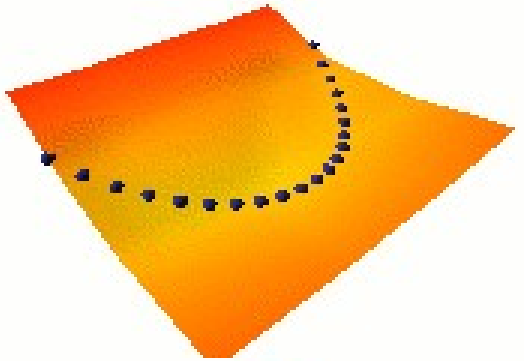

- Negative Log Likelihood

$$\mathcal{L}_{\text{nll}}(W, \mathcal{S}) = \frac{1}{P}\sum_{i=1}^{P}\left(E(W, Y^i, X^i) + \frac{1}{\beta}\log\int_{y\in\mathcal{Y}} e^{-\beta E(W,y,X^i)}\right).$$

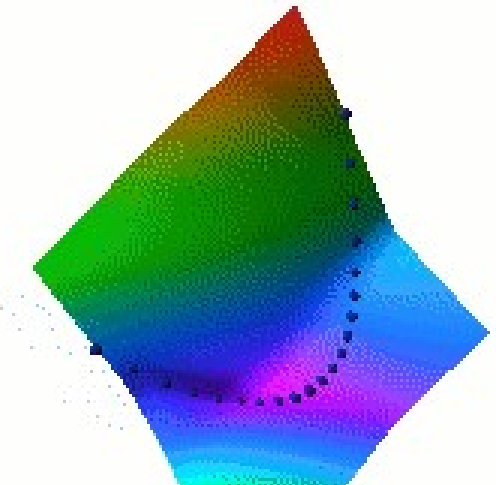

Loss Function to train Energy-Based Models

Good and bad loss functions

A tutorial on Energy-Based Learning [LeCun et al 2006]

Loss (equation #)	Formula	Margin
energy loss	$E(W, Y^i, X^i)$	none
perceptron	$E(W, Y^i, X^i) - \min_{Y \in \mathcal{Y}} E(W, Y, X^i)$	0
hinge	$\max\left(0, m + E(W, Y^i, X^i) - E(W, \bar{Y}^i, X^i)\right)$	m
log	$\log\left(1 + e^{E(W,Y^i,X^i) - E(W,\bar{Y}^i,X^i)}\right)$	> 0
LVQ2	$\min\left(M, \max(0, E(W, Y^i, X^i) - E(W, \bar{Y}^i, X^i)\right)$	0
MCE	$\left(1 + e^{-\left(E(W,Y^i,X^i) - E(W,\bar{Y}^i,X^i)\right)}\right)^{-1}$	> 0
square-square	$E(W, Y^i, X^i)^2 - \left(\max(0, m - E(W, \bar{Y}^i, X^i))\right)^2$	m
square-exp	$E(W, Y^i, X^i)^2 + \beta e^{-E(W,\bar{Y}^i,X^i)}$	> 0
NLL/MMI	$E(W, Y^i, X^i) + \frac{1}{\beta} \log \int_{y \in \mathcal{Y}} e^{-\beta E(W,y,X^i)}$	> 0
MEE	$1 - e^{-\beta E(W,Y^i,X^i)} / \int_{y \in \mathcal{Y}} e^{-\beta E(W,y,X^i)}$	> 0

"Shallow" Structured Prediction

Energy function is linear in the parameters

$$E(X,Y,Z) = \sum_i W_i^T h_i(X,Y,Z)$$

with the NLL Loss :

- **Conditional Random Field** [Lafferty, McCallum, Pereira 2001]

with Hinge Loss:

- **Max Margin Markov Nets** and **Latent SVM** [Taskar, Altun, Hofmann...]

with Perceptron Loss

- **Structured Perceptron** [Collins...]

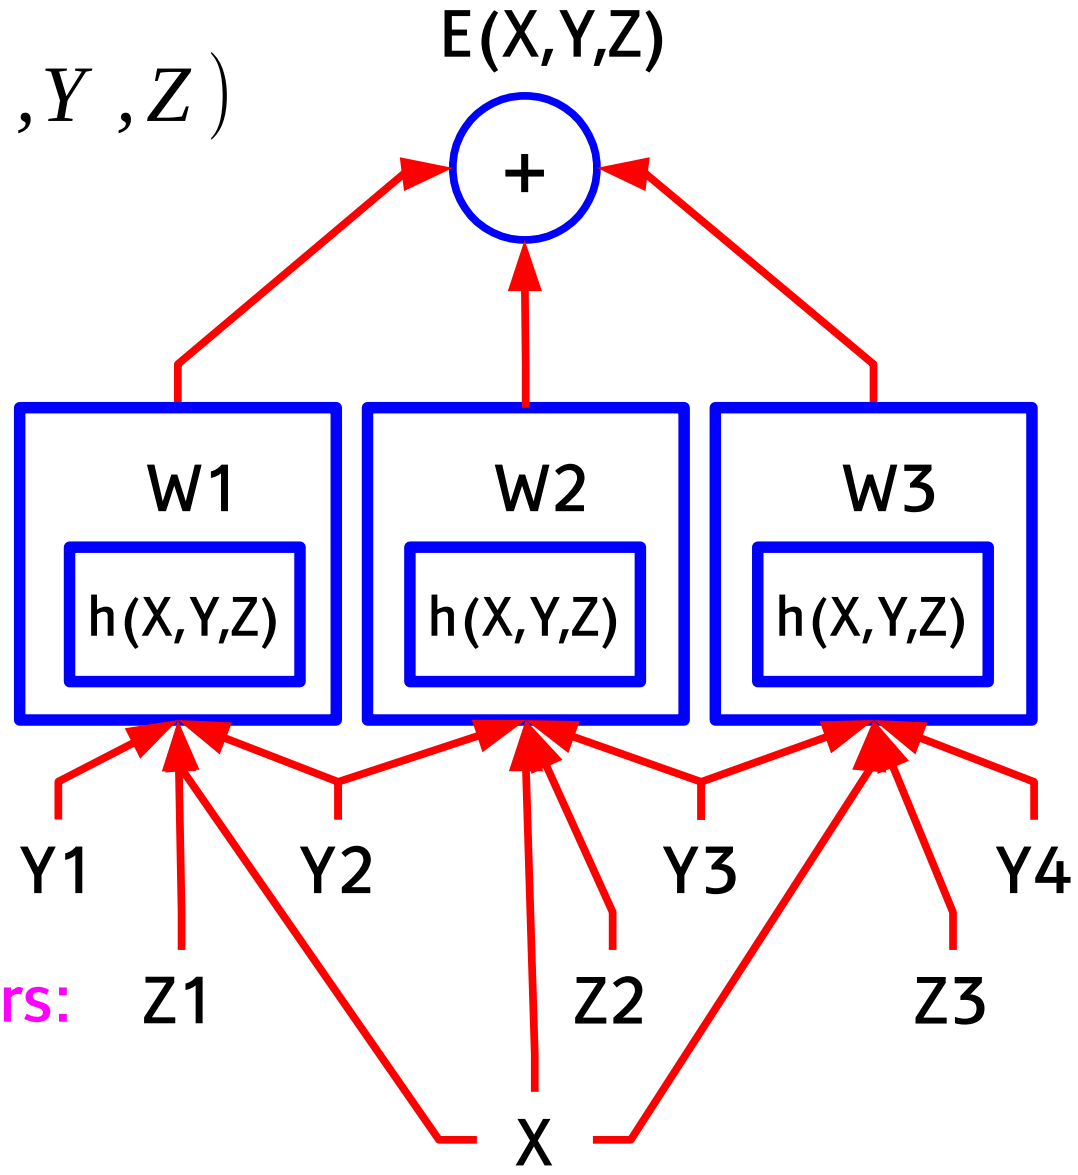

Energy function is linear in the parameters

$$E(X,Y,Z)=\sum_i g_i(X,Y,Z,W_i)$$

Graph Transformer Networks

- [LeCun, Bottou, Bengio,Haffner 97,98]
- NLL loss
- Perceptron loss

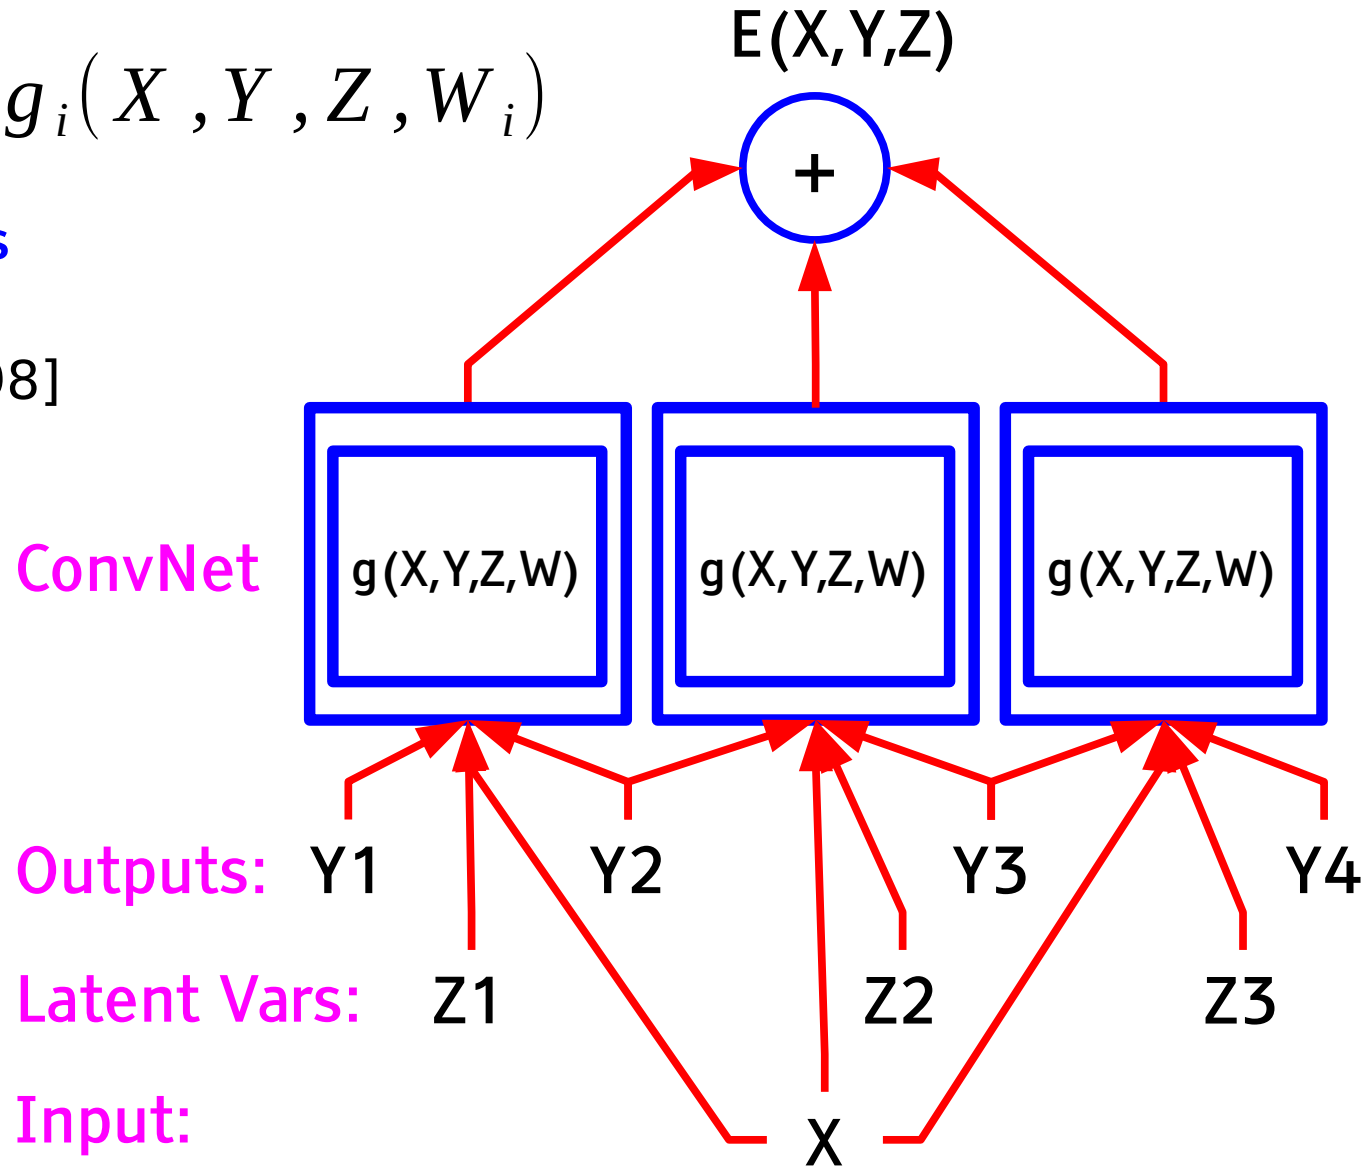

Integrating Deep Learning and Structured Prediction

- **Deep Net + Graphical Model = Factor Graph**
  - Energy function is a sum of factors
  - Factors can embed deep architectures
  - X: observed variables (inputs)
  - Z: never observed (latent variables)
  - Y: observed on training set (output variables)
- **Inference is energy minimization (MAP) or free energy minimization (marginalization) over Z and Y given an X**
  - F(X,Y) = MIN_z E(X,Y,Z)
  - F(X,Y) = -log[SUM_z exp(-E(X,Y,Z))]

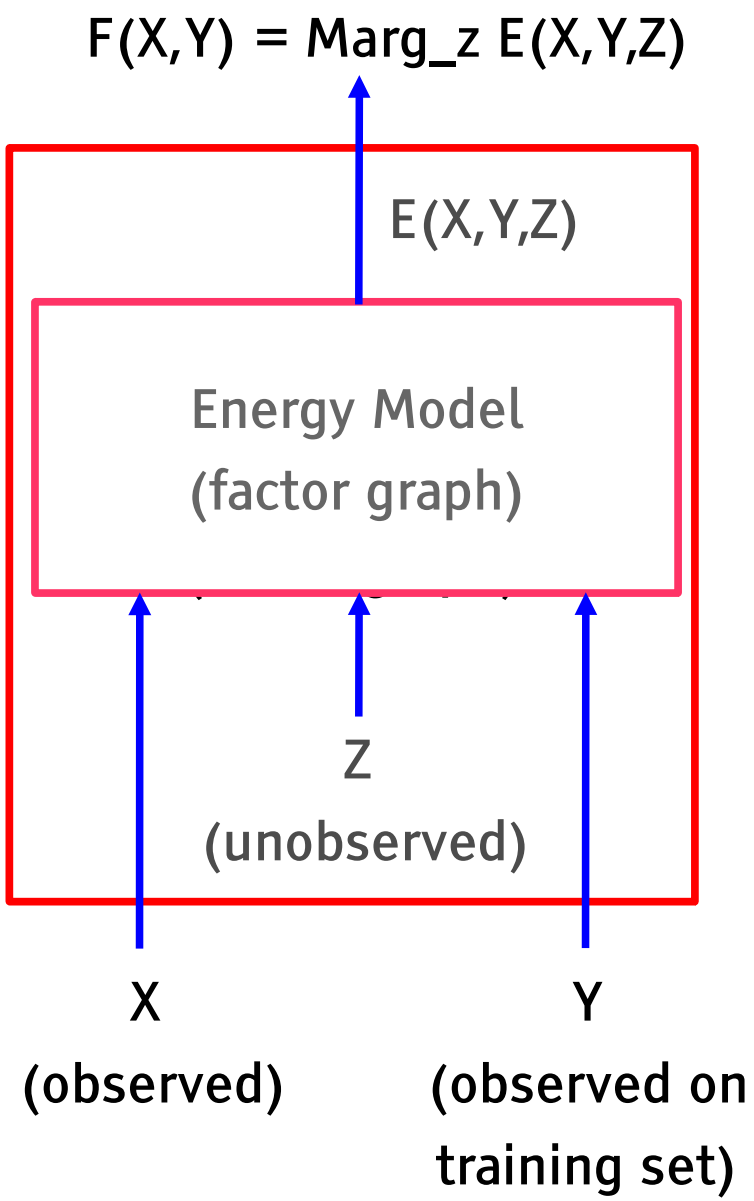

End-to-End Learning – Word-Level Training

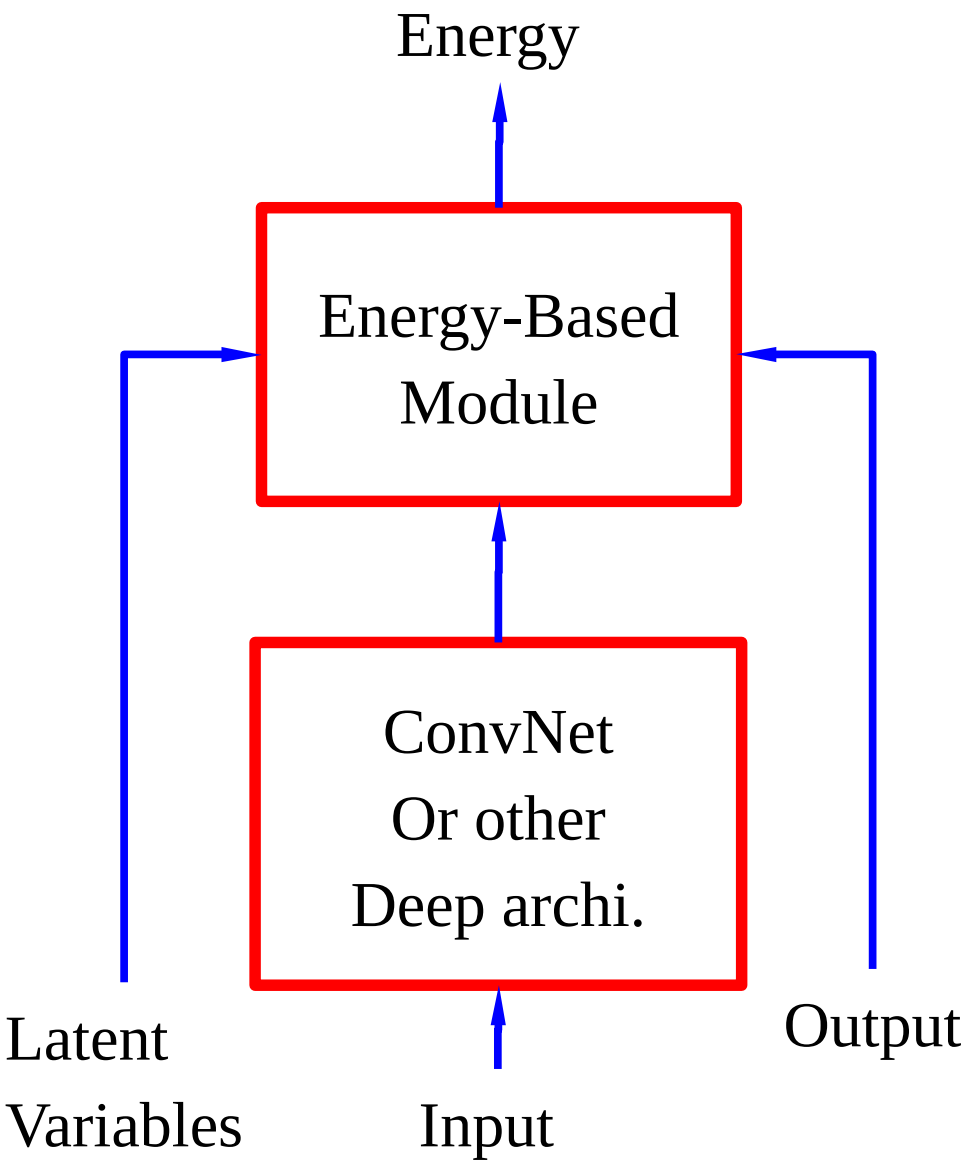

Making every single module in the system trainable.

Every module is trained simultaneously so as to optimize a global loss function.

Includes the feature extractor, the recognizer, and the contextual post-processor (graphical model)

Problem: back-propagating gradients through the graphical model.

Deep Structured Prediction for Speech Recognition (1991)

Spoken word recognition with trainable elastic templates and trainable feature extraction [Driancourt&Bottou 1991, Bottou 1991, Driancourt 1994]

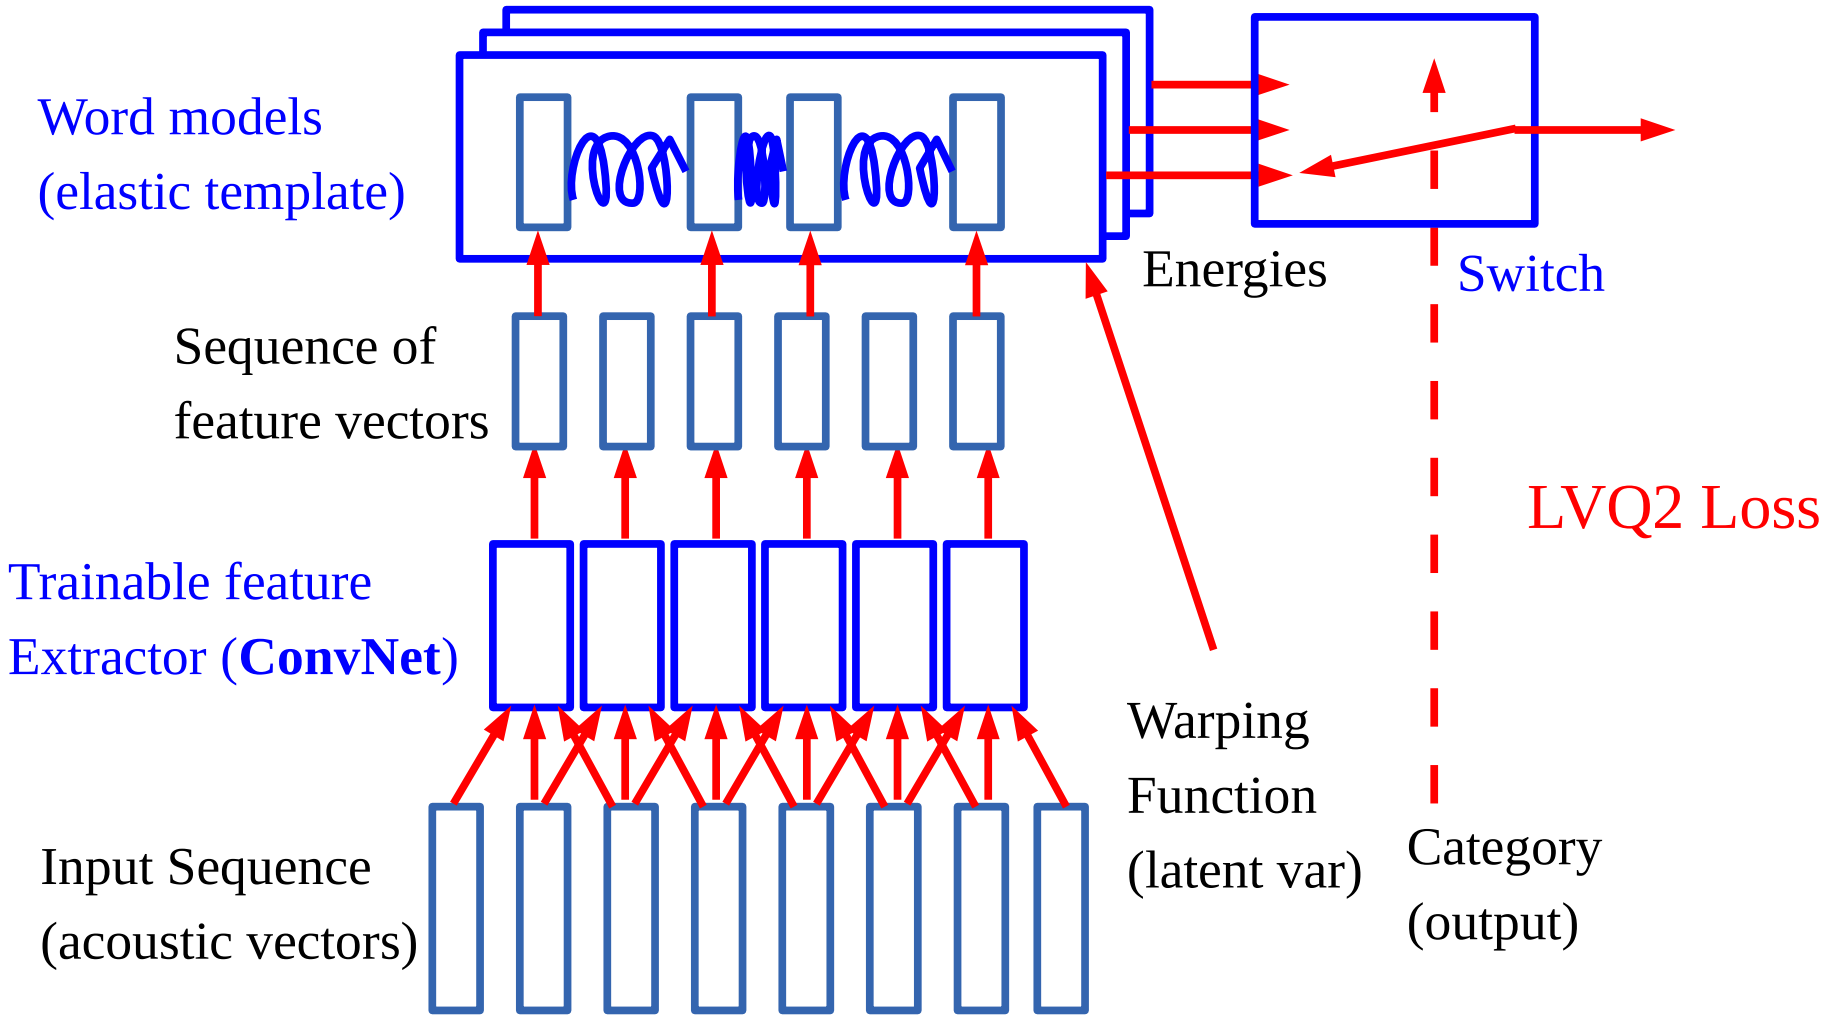

Alignment through Dynamic Time Warping

Spoken word recognition with trainable elastic templates and trainable feature extraction

[Driancourt&Bottou 1991, Bottou 1991, Driancourt 1994]

Elastic matching using dynamic time warping (Viterbi algorithm on a trellis).

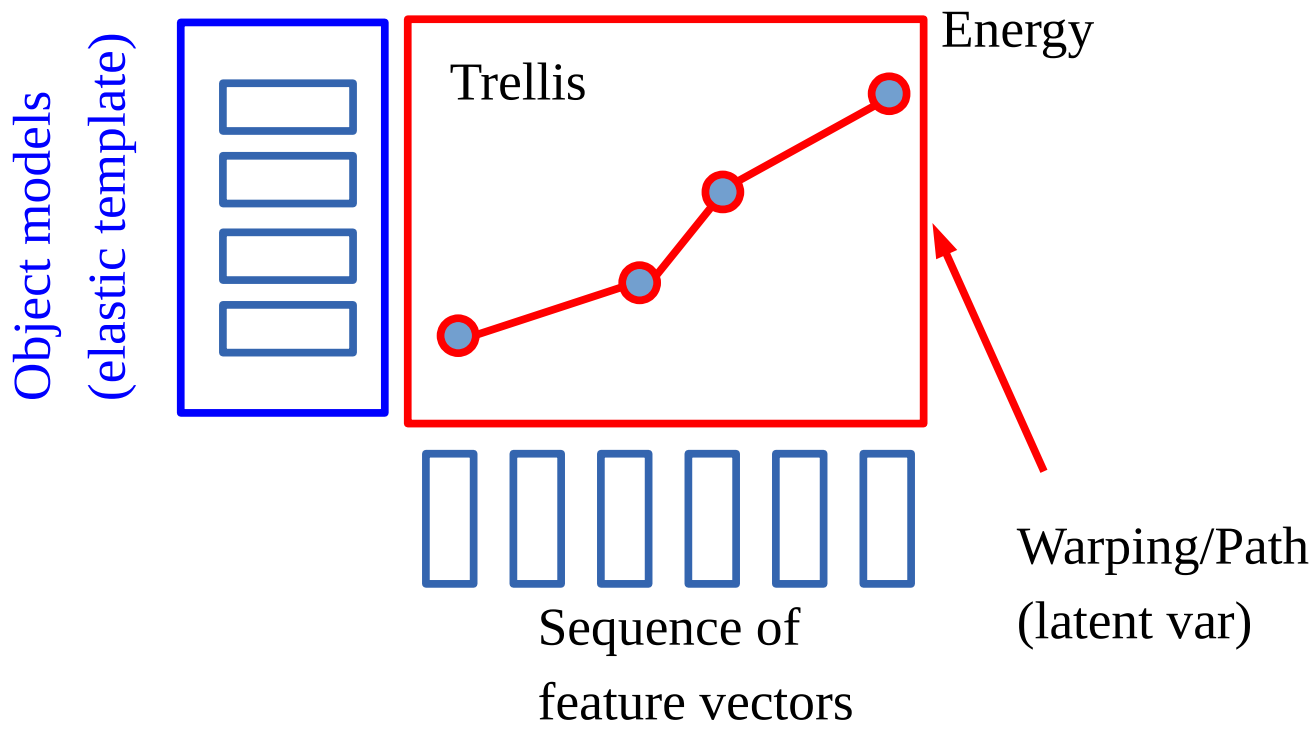

Deep Factors / Deep Graph: ASR with TDNN/DTW

Trainable Automatic Speech Recognition system with convolutional nets (TDNN) and dynamic time warping (DTW)

Training the feature extractor as part of the whole process.

with the LVQ2 Loss :

- Driancourt and Bottou's speech recognizer (1991)

with NLL:

- Bengio's speech recognizer (1992)
- Haffner's speech recognizer (1993)

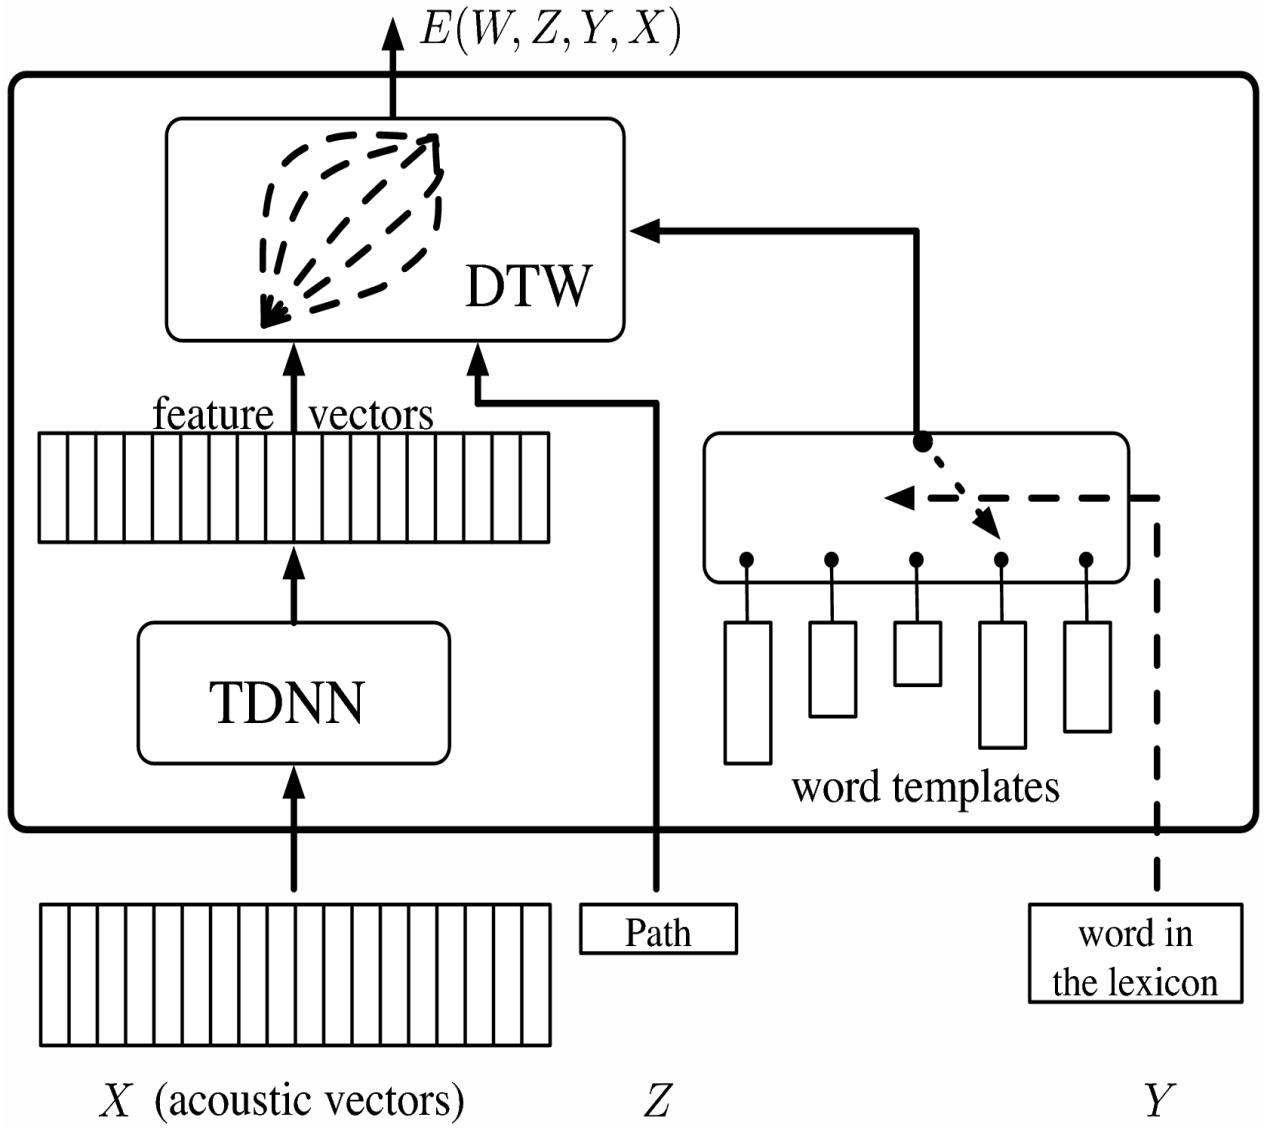

Using Graphs instead of Vectors or Arrays.

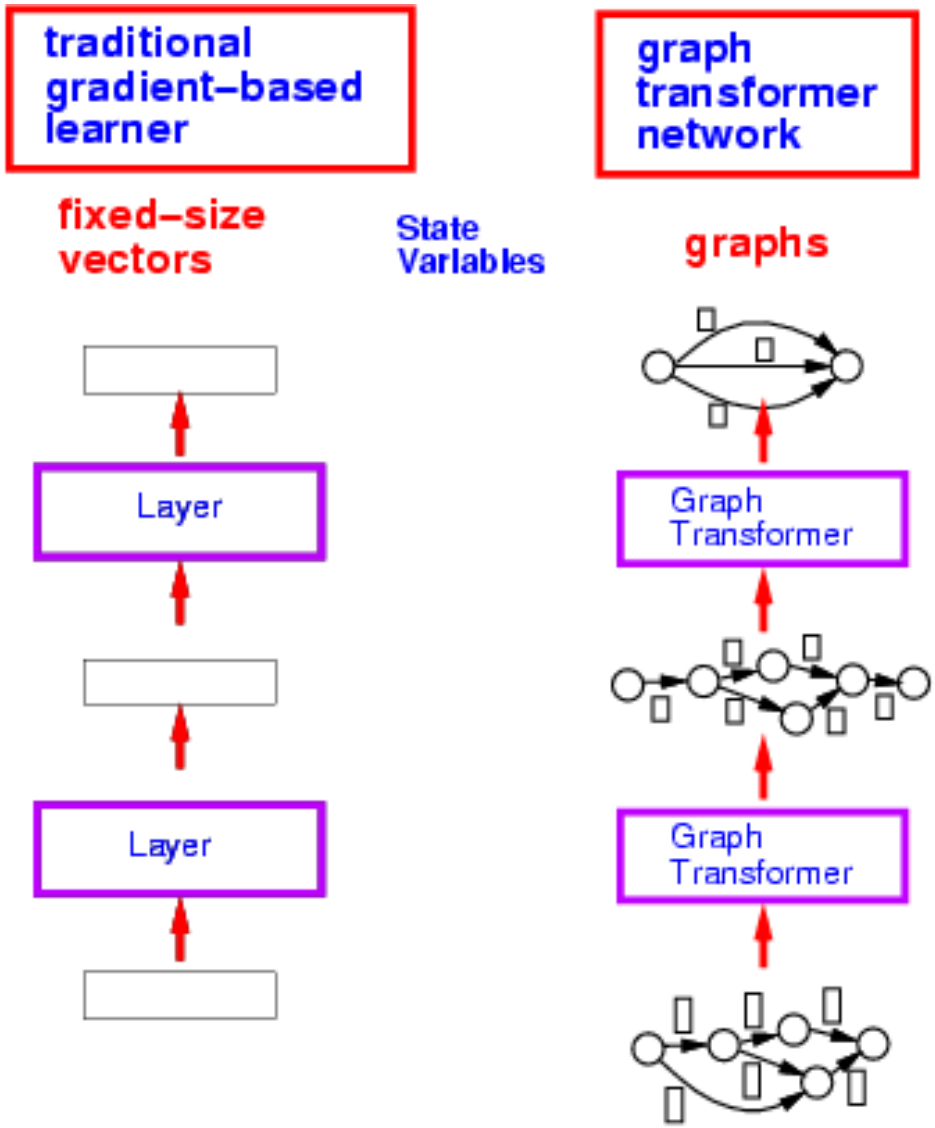

Whereas traditional learning machines manipulate fixed-size vectors, Graph Transformer Networks manipulate graphs.

Graph Transformer Networks

Variables:

- X: input image
- Z: path in the interpretation graph/segmentation
- Y: sequence of labels on a path

Loss function: computing the energy of the desired answer:

$$E(W, Y, X)$$

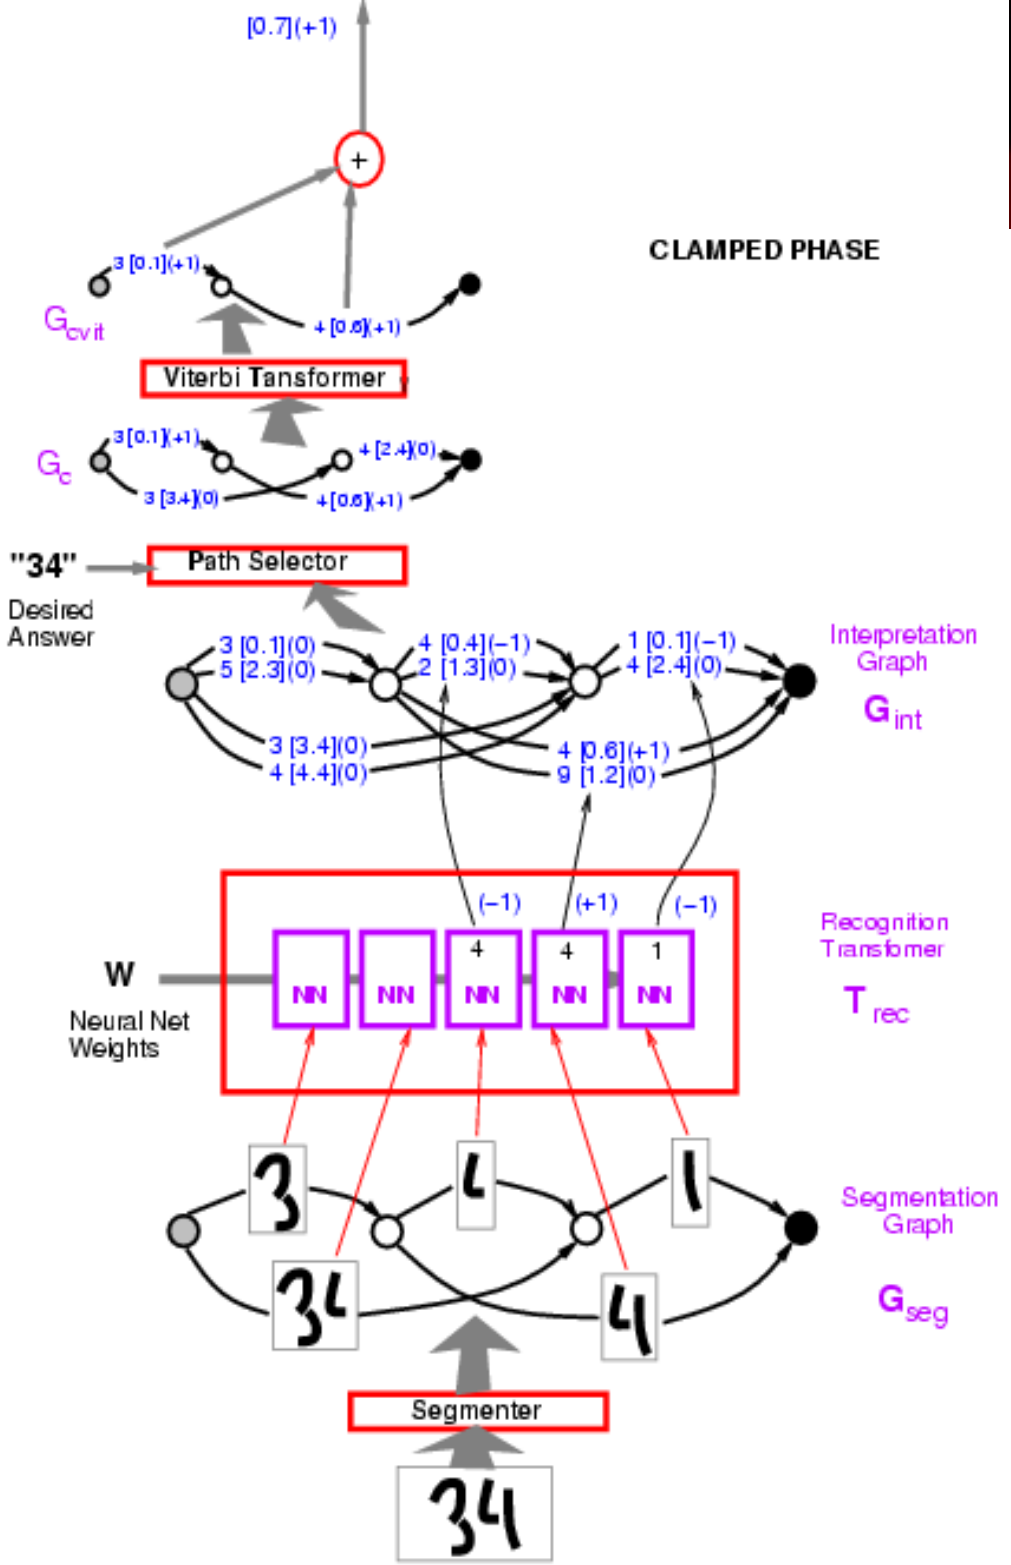

Graph Transformer Networks

Variables:

- X: input image
- Z: path in the interpretation graph/segmentation
- Y: sequence of labels on a path

Loss function: computing the constrastive term:

$$E(W, \check{Y}, X)$$

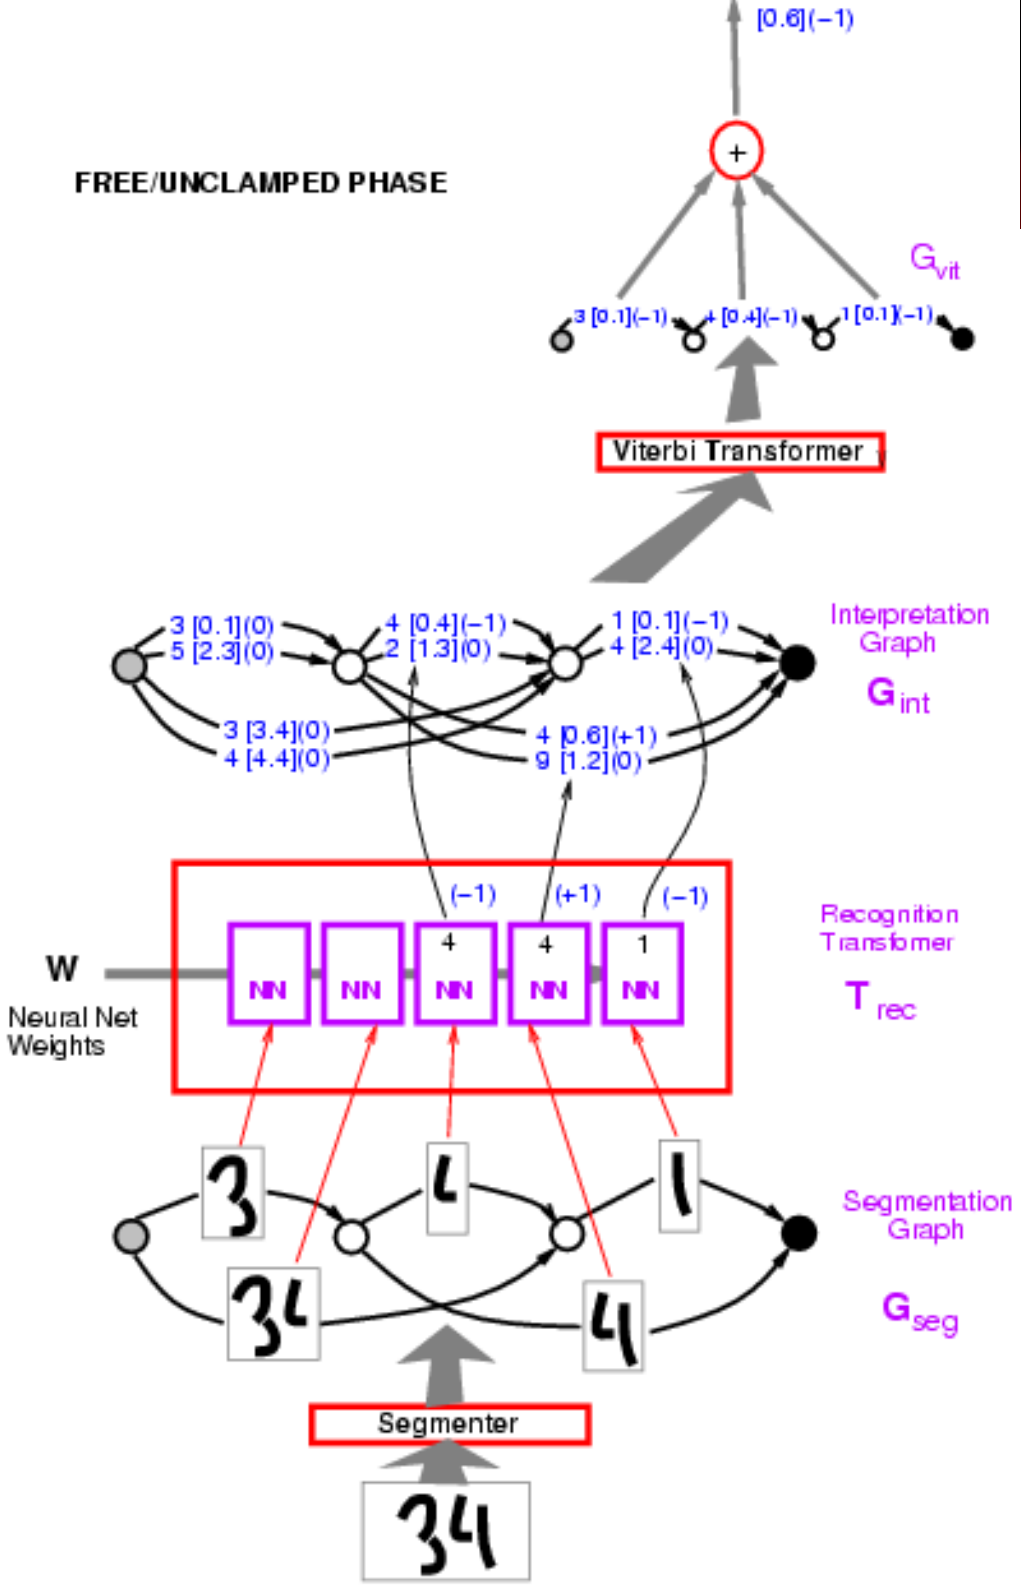

Graph Transformer Networks

Example: Perceptron loss

Loss = Energy of desired answer – Energy of best answer.

- (no margin)

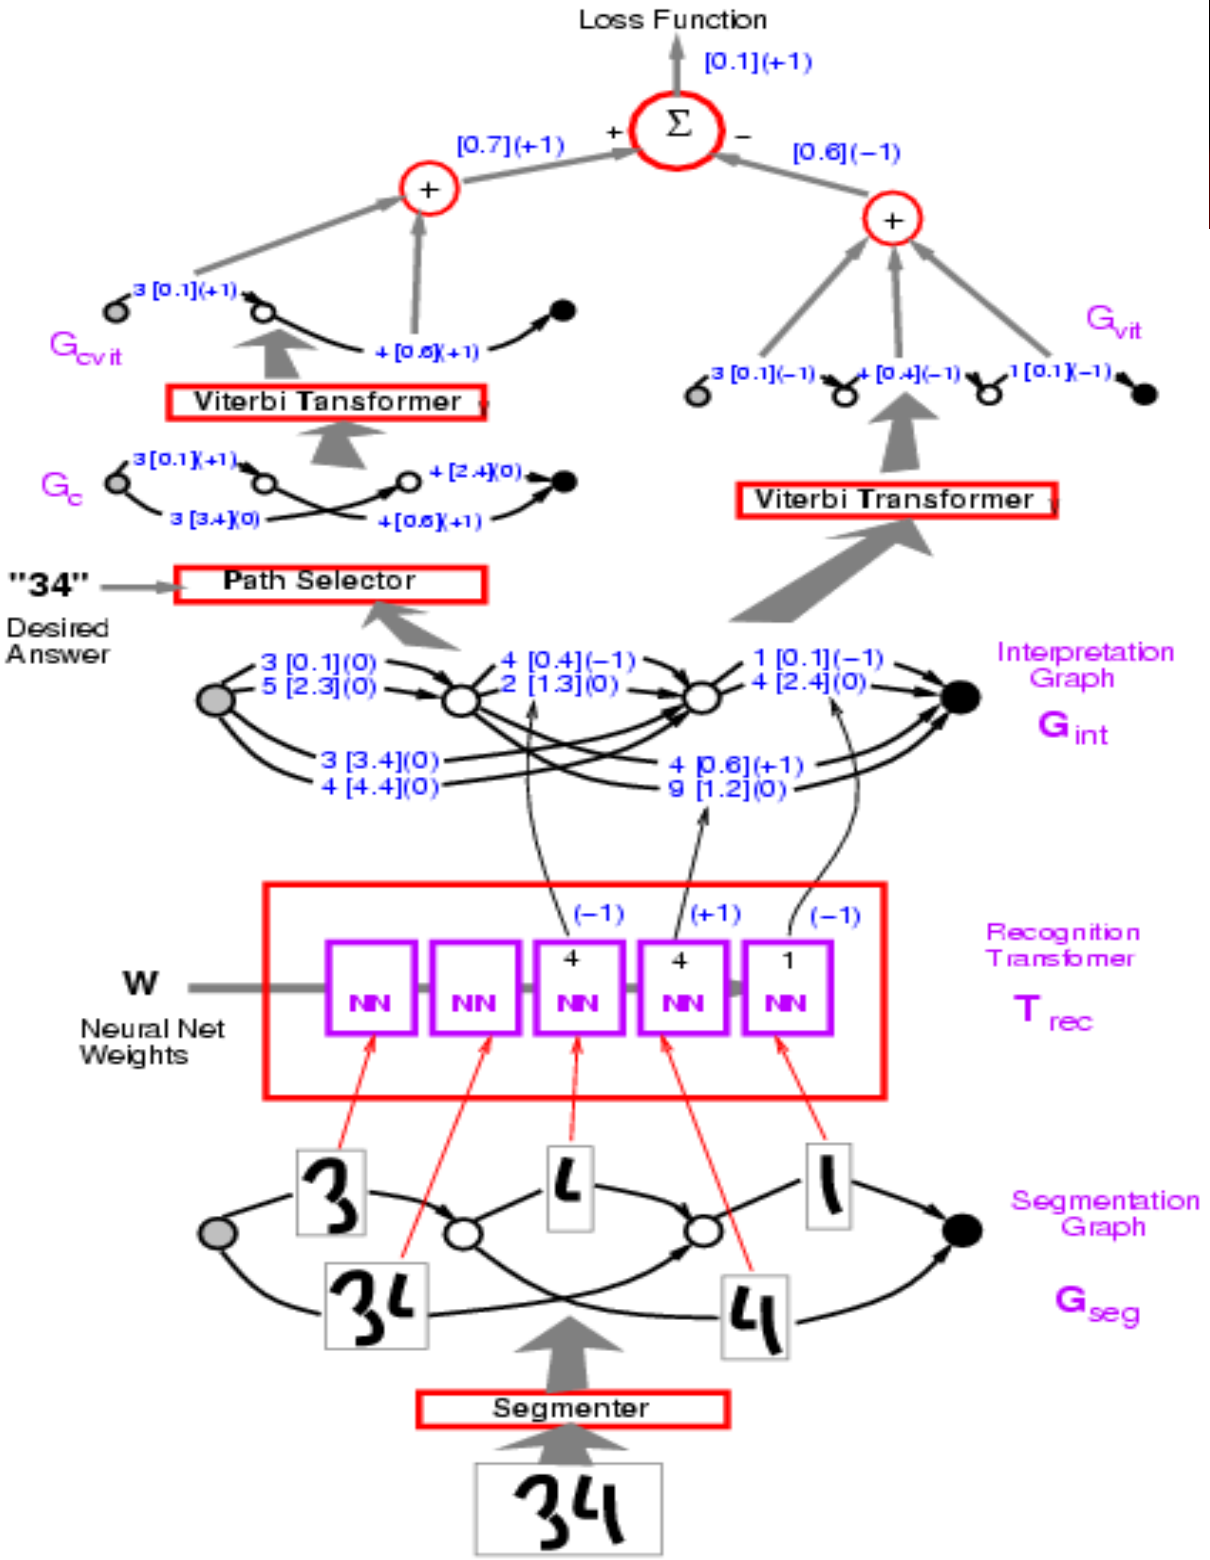

Integrating Deep Learning and Structured Prediction

Y LeCun

- **Structured prediction: when the output is structured: string, graph.....**
- **Integrating deep learning and structured prediction is an old idea**
  - In fact, it predates structured prediction [LeCun, Bottou, Bengio, Haffner 1998]
- **Globally-trained convolutional-net + graphical models for handwriting recognition**
  - trained discriminatively at the word level
  - Loss identical to CRF and structured perceptron
  - Compositional movable parts model

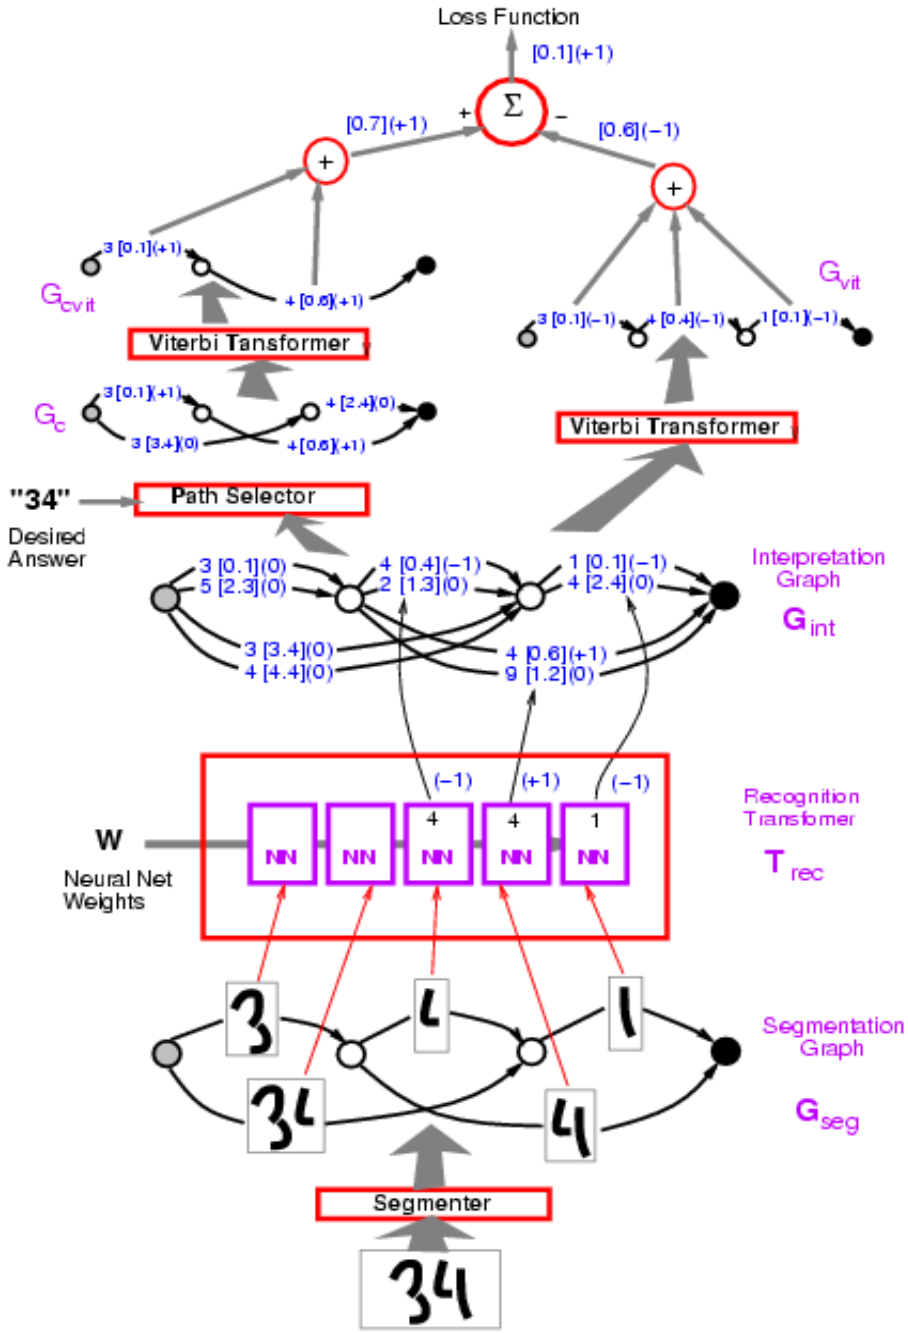

Global (word-level) Training Helps

Pen-based handwriting recognition (for tablet computer)

- [Bengio&LeCun 1995]

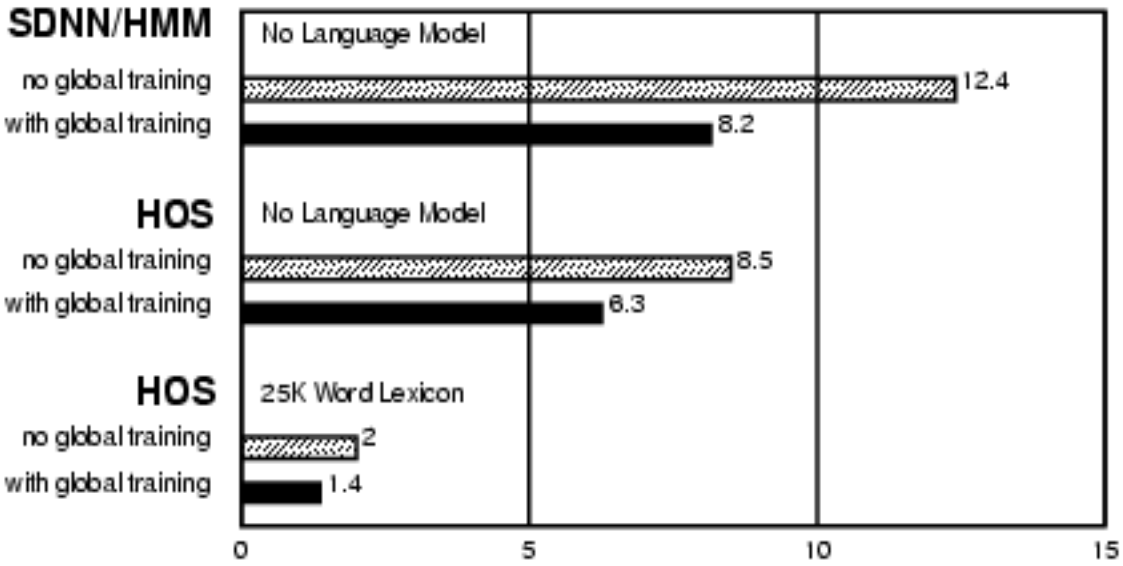

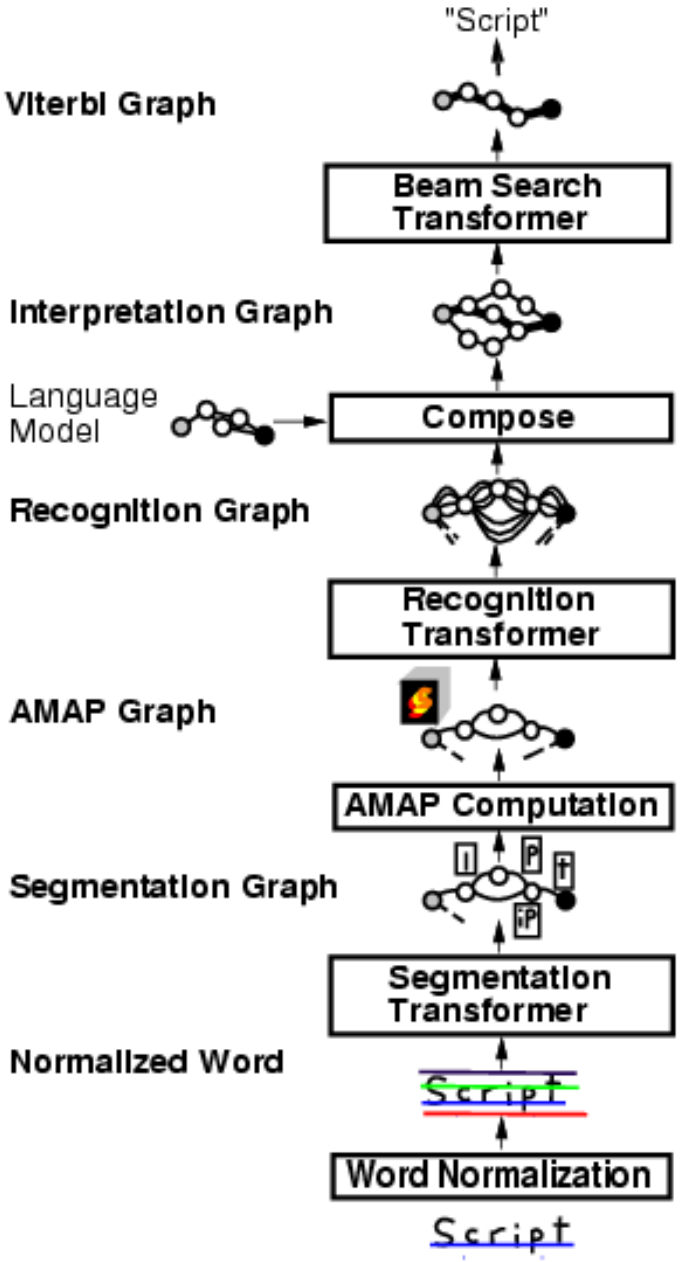

Graph Composition, Transducers.

The composition of two graphs can be computed, the same way the dot product between two vectors can be computed.

General theory: semi-ring algebra on weighted finite-state transducers and acceptors.

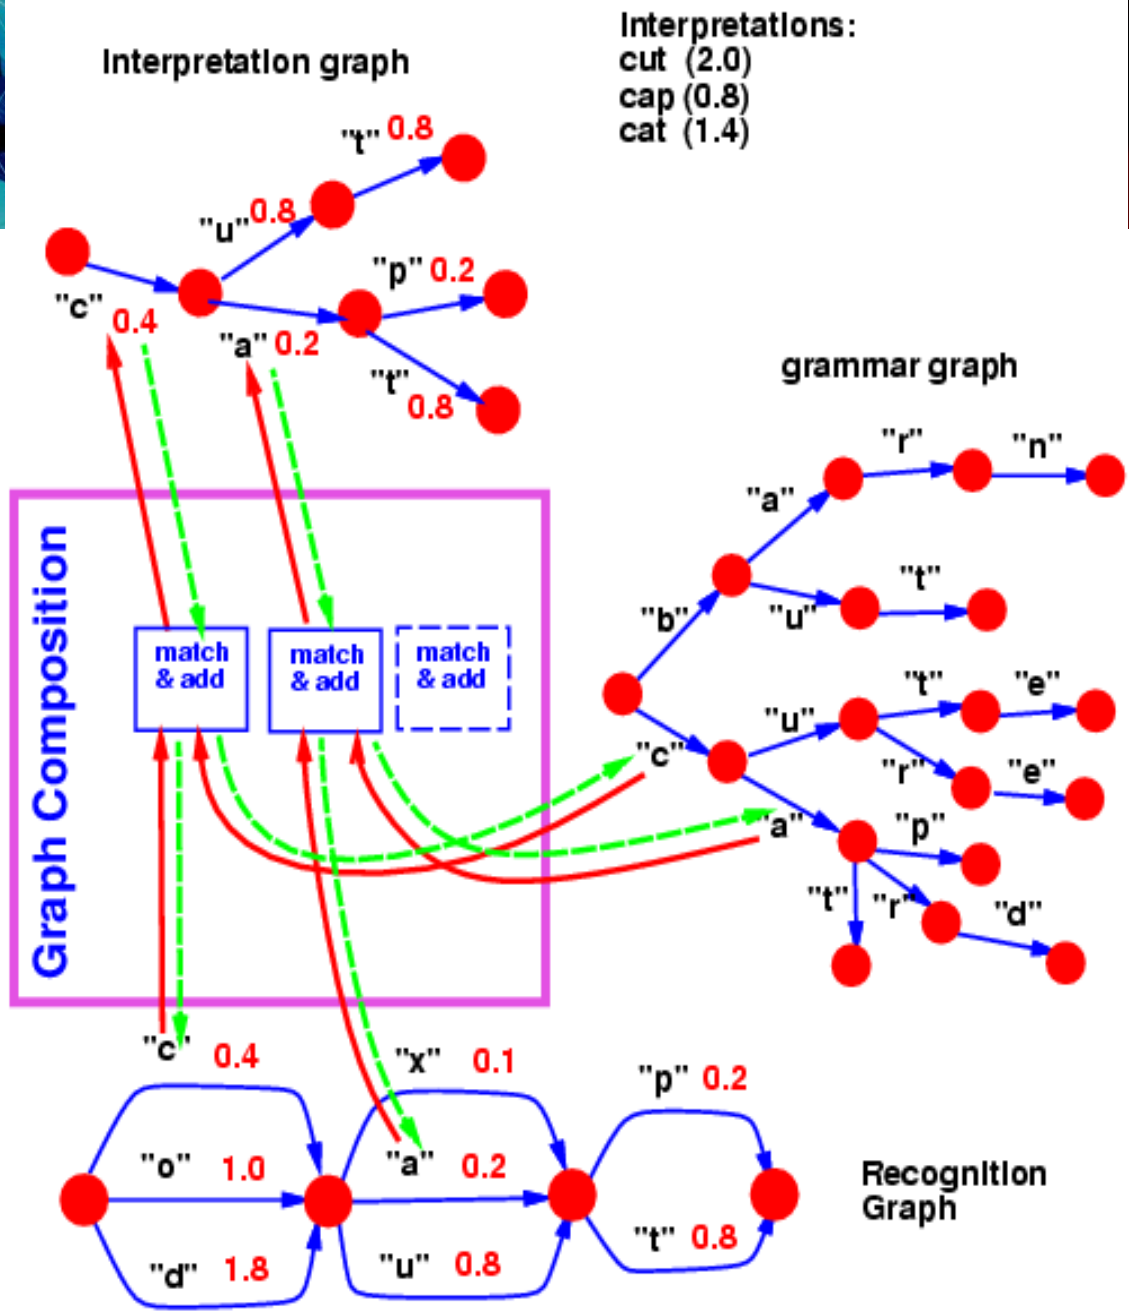

Check Reader

Graph transformer network trained to read check amounts.

Trained globally with Negative-Log-Likelihood loss.

50% percent correct, 49% reject, 1% error (detectable later in the process.

Fielded in 1996, used in many banks in the US and Europe.

Processes an estimated 10% to 20% of all the checks written in the US.

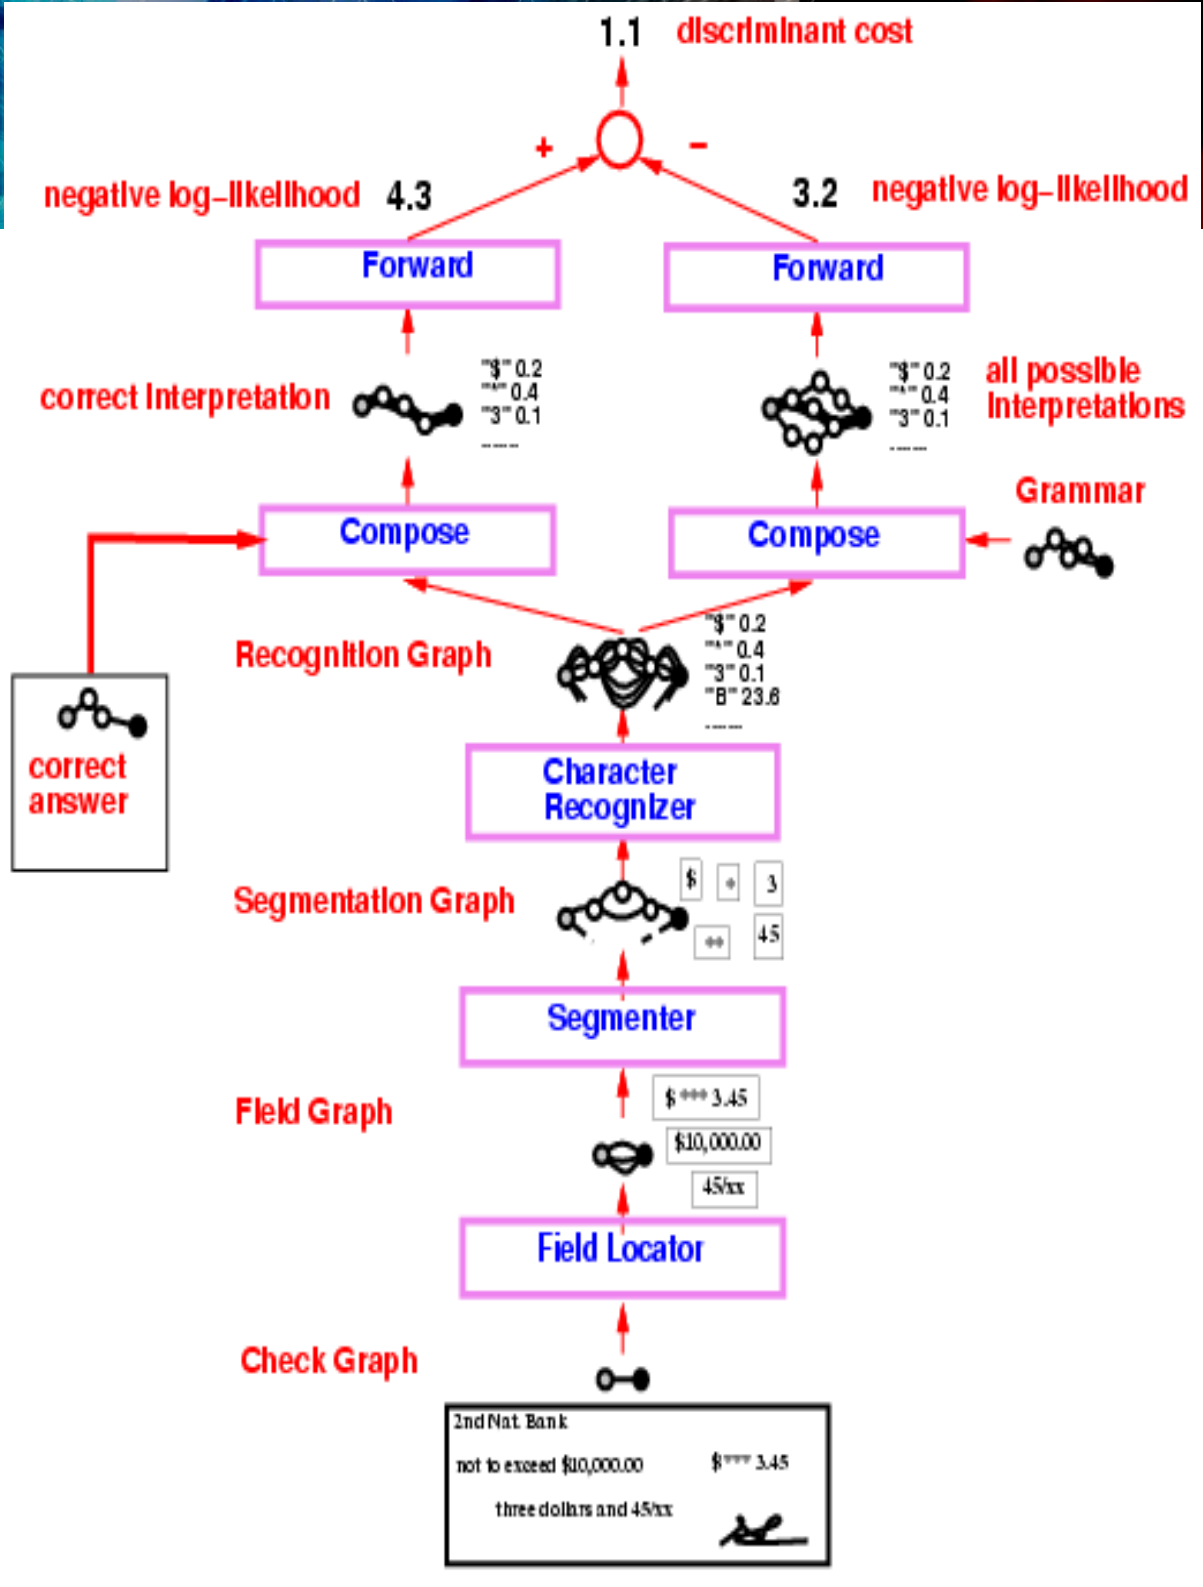

Deep Structured Predictions for Speech and Handwriting

Trainable Speech/Handwriting Recognition systems that integrate Neural Nets (or other "deep" classifiers) with dynamic time warping, Hidden Markov Models, or other graph-based hypothesis representations

Word-level global discriminative training with GMM:

With Minimum Empirical Error loss

- Ljolje and Rabiner (1990)

With MCE

- Juang et al. (1997)

Word-level global discriminative training with ConvNets:

with the LVQ2 Loss :

- Driancourt and Bottou's speech recognizer (1991)

with Neg Log Likelihood (aka MMI):

- Bengio (1992), Haffner (1993), Bourlard (1994)

CRF-like Late normalization

- un-normalized HMM
- Bottou pointed out the **label bias problem** (1991)
- Denker and Burges proposed a solution (1995)
- Implemented in (LeCun et al 1998)

Deep Learning in
Natural Language Processing
Research at Facebook

TagSpace: hash tag prediction [Adams et al. 2014]

We want to learn semantic features for post content.

Hashtags are a cheap, plentiful labeling of text provided by authors, similar to classical NLP task labels:

disambiguation (chips #futurism vs. chips #junkfood); identification of named entities (#sf49ers); sentiment (#dislike); and topic annotation (#yoga).

We train a neural network using hashtags as supervisio

training on 5.5 billion words and 100,000 possible hashtags.

The learned features are useful for other tasks, e.g. food r

"crazy commute this am"	"marriage equality now"
#snow, #snowstorm, #pax, #tubestrike #snowpocalypse,	#samelove, #lgbt, #equality, #equalrights,

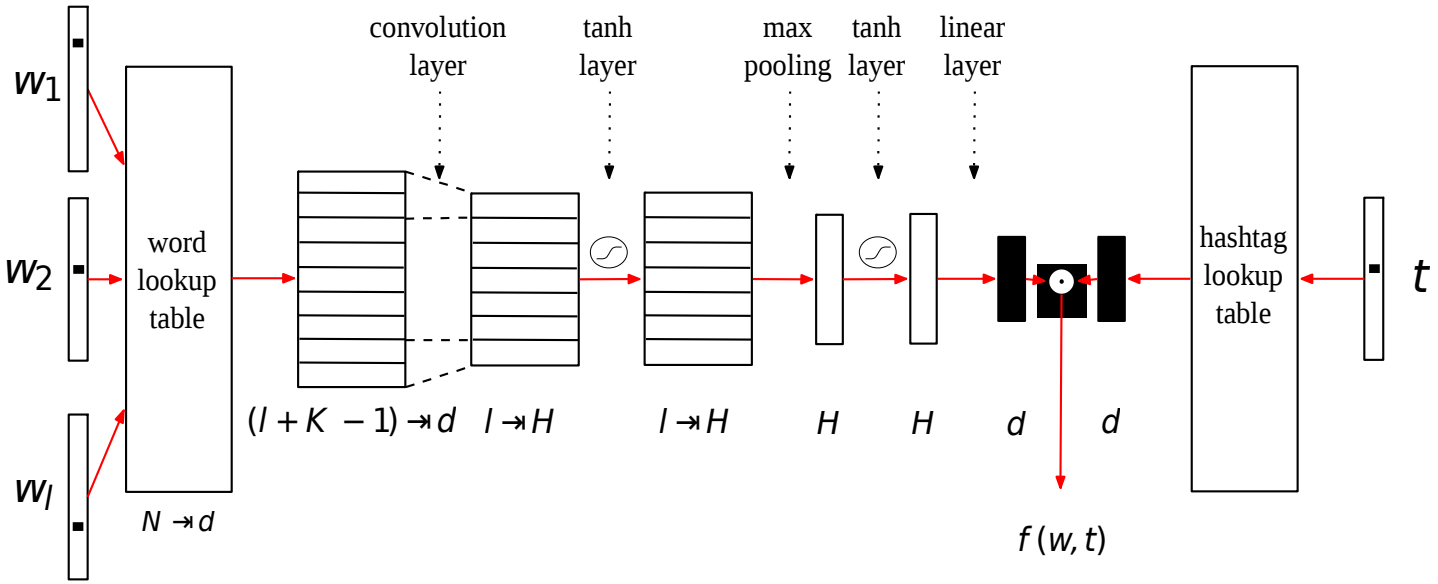

"kevin spacey is a super hottie"
#hoc, #mcm, #hocseason2, #season2

Predicting Posts and Comments with Recurrent Nets

- happy birthday to you ! big 17 ! !
- ready to go on the swim alone ! :)
- the fbi devil hits my folder but really really crazy . its not my bro ? ! ?
- happy birthday UNK ! ! i hope you have a great day . .
- i have an urgent candy crush exam . she UNK some awesome service on plus june 7th \u0040 UNK camp 2014 center , 8pm central to go watch it now ? it 's not hurt . everybody know it 's been a long journey . watch it UNK the UNK islands of earth . y'all stay safe UNK UNK love ya honey ! thank you so much for all the flowers and nurses for me . i will bounce down until next temple but please stop following to helping my friend as they have no choice like this . \ women are the ones i used to ?
- so proud of jack williams . thanks to all the teachers who have served to send respect to every faithfulness and generations to us all .
- # UNK UNK insecurity rose so high that our UNK ladies no longer have even safe days !
- hey boss ! i was the change of UNK . . we have a party together on friday UNK ba ang mga bum ?

Question-Answering System

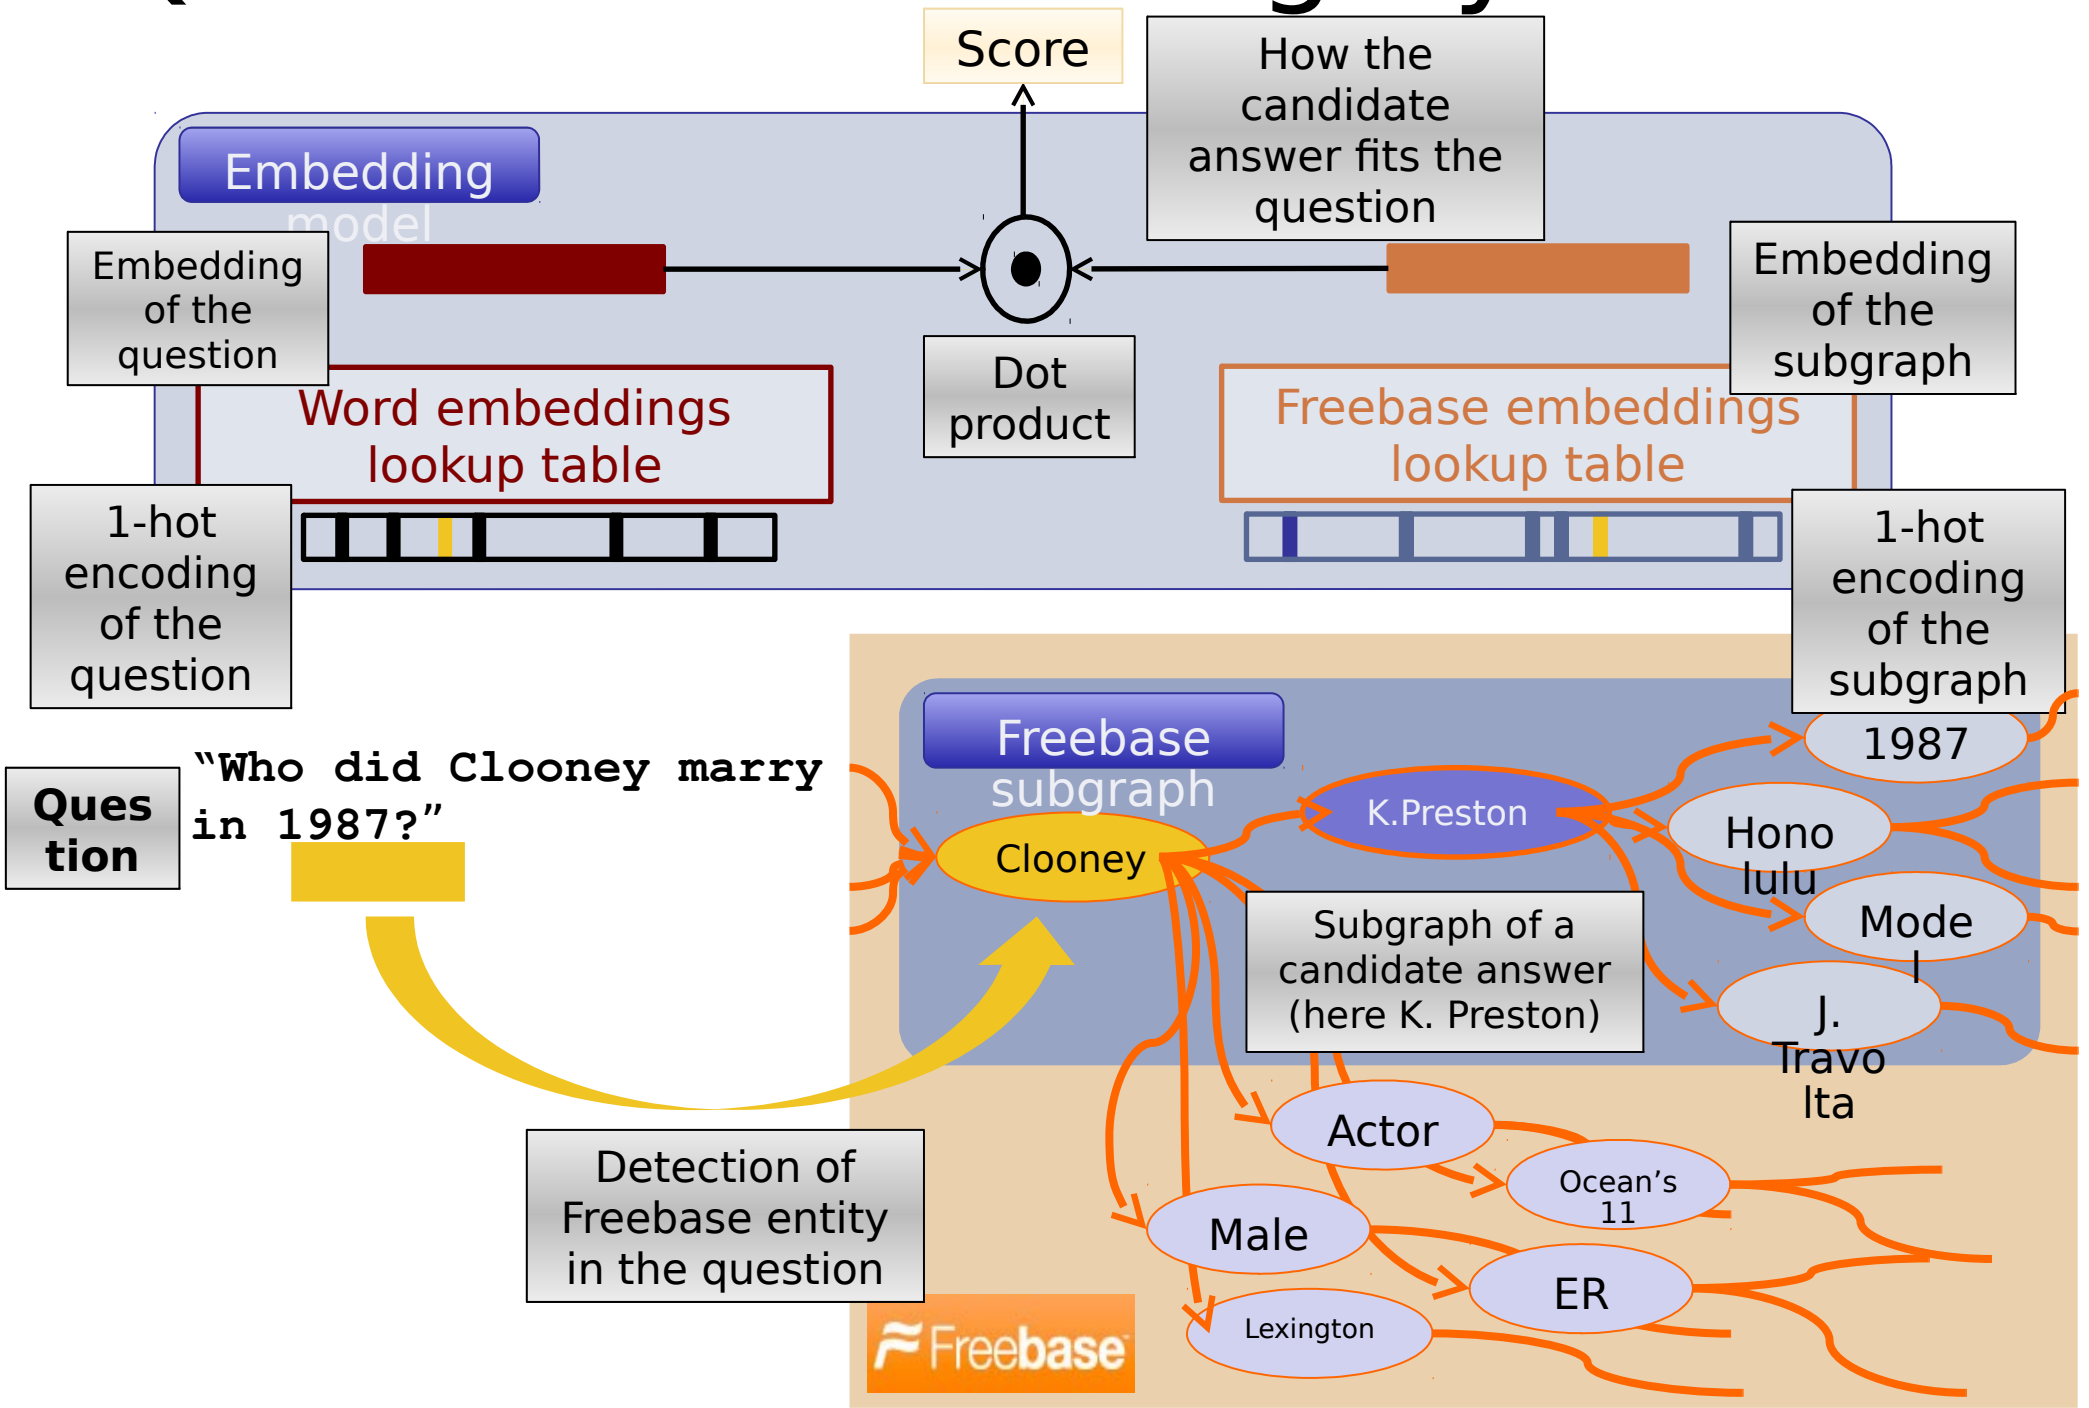

Question-Answering System

what are bigos?
["stew"] ["stew"]
what are dallas cowboys colors?
["navy_blue", "royal_blue", "blue", "white", "silver"] ["blue", "navy_blue", "white", "royal_blue", "silver"]
how is egyptian money called?
["egyptian_pound"] ["egyptian_pound"]
what are fun things to do in sacramento ca?
["sacramento_zoo"] ["raging_waters_sacramento", "sutter_s_fort", "b_street_theatre", "sacramento_zoo", "california_state_capitol_museum", ...]
how are john terry's children called?
["georgie_john_terry", "summer_rose_terry"] ["georgie_john_terry", "summer_rose_terry"]
what are the major languages spoken in greece?
["greek_language", "albanian_language"] ["greek_language", "albanian_language"]
what was laura ingalls wilder famous for?
["writer", "author"] ["writer", "journalist", "teacher", "author"]

Question-Answering System

who plays sheldon cooper mother on the big bang theory?
["jim_parsons"] ["jim_parsons"]
who does peyton manning play football for?
["denver_broncos"] ["indianapolis_colts", "denver_broncos"]
who did vladimir lenin marry?
["nadezhda_krupskaya"] ["nadezhda_krupskaya"]
where was teddy roosevelt's house?
["new_york_city"] ["manhattan"]
who developed the tcp ip reference model?
["vint_cerf", "robert_e._kahn"] ["computer_scientist", "engineer”]

Memory Network [Weston, Chopra, Bordes 2014]

- **Add a short-term memory to a network** http://arxiv.org/abs/1410.3916

Joe went to the garden then Fred picked up the milk; Joe moved to the bathroom and Fred dropped the milk, and then Dan moved to the living room.

- *Where is Dan?* A: living room I believe
- *Where is Joe?* A: the bathroom

Fred moved to the bedroom and Joe went to the kitchen then Joe took the milk there and Dan journeyed to the bedroom; Joe discarded the milk.

- *Where is the milk now?* A: the milk is in the kitchen
- *Where is Dan now?* A: I think he is in the bedroom

Joe took the milk there, after that Mike travelled to the office, then Joe went to the living room, next Dan went back to the kitchen and Joe travelled to the office.

- *Where is Joe now?* A: I think Joe is in the office

Memory Network [Weston, Chopra, Bordes 2014]

- Add a short-term memory to a network

http://arxiv.org/abs/1410.3916

I: (input feature map) – converts the incoming input to the internal feature representation.
G: (generalization) – updates old memories given the new input.
O: (output feature map) – produces a new output (in the feature representation space), given the new input and the current memory.
R: (response) – converts the output into the response format desired. For example, a textual response or an action.

Method	F1
(Fader et al., 2013) [4]	0.54
(Bordes et al., 2014) [3]	0.73
MemNN	0.71
MemNN (with BoW features)	0.79

Results on Question Answering Task

Bilbo travelled to the cave.
Gollum dropped the ring there.
Bilbo took the ring.
Bilbo went back to the Shire.
Bilbo left the ring there.
Frodo got the ring.
Frodo journeyed to Mount-Doom.
Frodo dropped the ring there.
Sauron died.
Frodo went back to the Shire.
Bilbo travelled to the Grey-havens.
The End.
Where is the ring? A: Mount-Doom
Where is Bilbo now? A: Grey-havens
Where is Frodo now? A: Shire

Fig. 2. An example story with questions correctly answered by a MemNN. The MemNN was trained on the simulation described in Section 4.2 and had never seen many of these words before, e.g. Bilbo, Frodo and Gollum.

Unsupervised Learning

Energy-Based Unsupervised Learning

- Learning an energy function (or contrast function) that takes
  - Low values on the data manifold
  - Higher values everywhere else

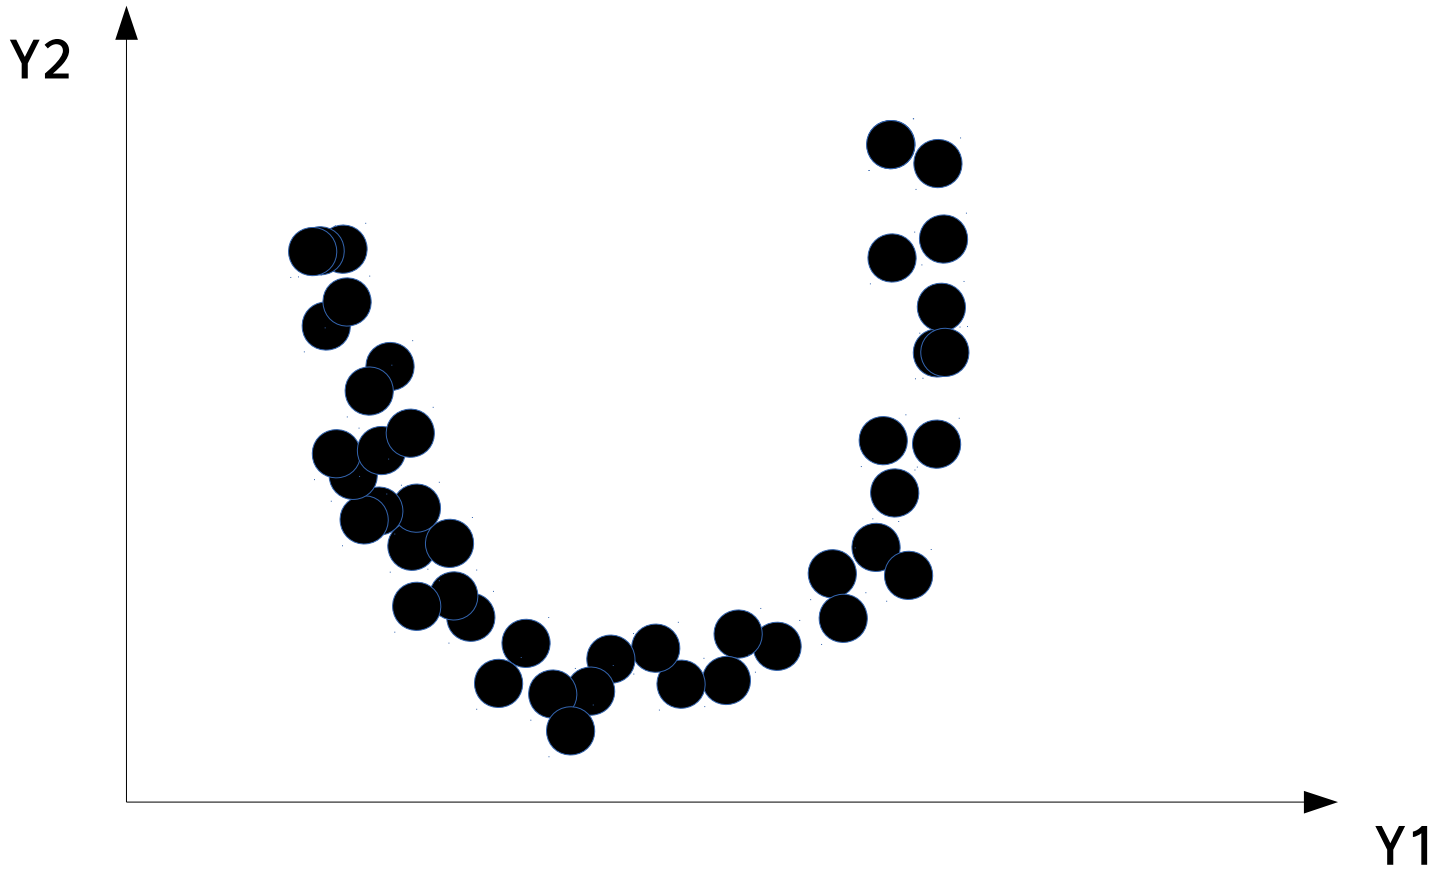

Capturing Dependencies Between Variables with an Energy Function

- The energy surface is a "contrast function" that takes low values on the data manifold, and higher values everywhere else
  - Special case: energy = negative log density
  - Example: the samples live in the manifold

$$Y_2 = (Y_1)^2$$

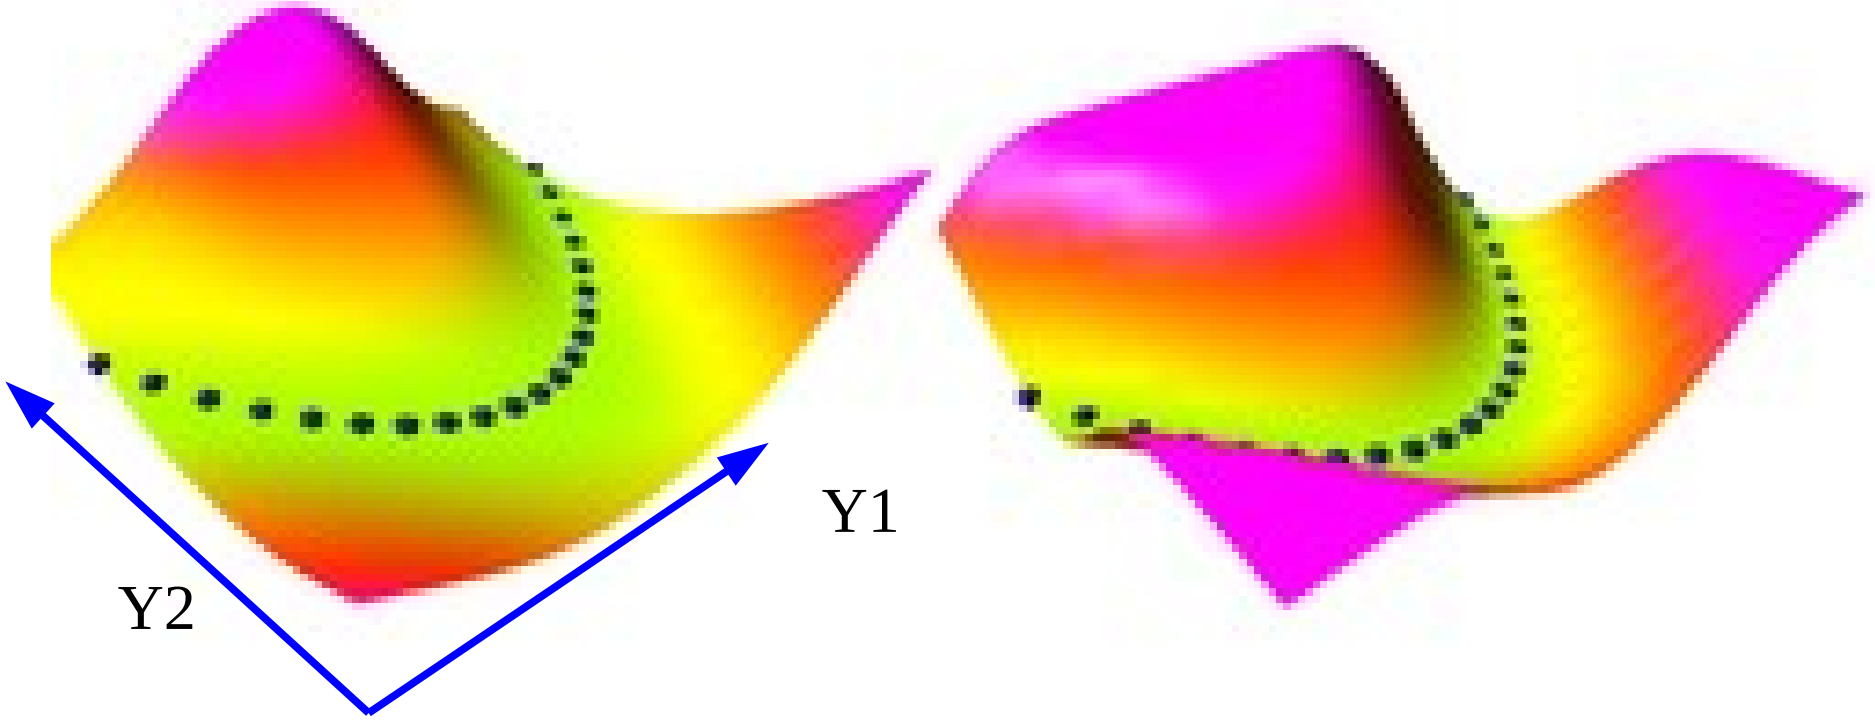

Transforming Energies into Probabilities (if necessary)

- The energy can be interpreted as an unnormalized negative log density
- Gibbs distribution: Probability proportional to exp(-energy)
  - Beta parameter is akin to an inverse temperature
- Don't compute probabilities unless you absolutely have to
  - Because the denominator is often intractable

$$P(Y|W) = \frac{e^{-\beta E(Y,W)}}{\int_y e^{-\beta E(y,W)}}$$

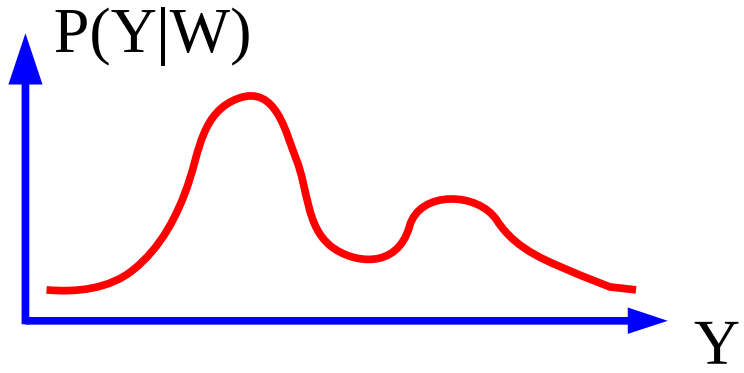

$$E(Y,W) \propto -\log P(Y|W)$$

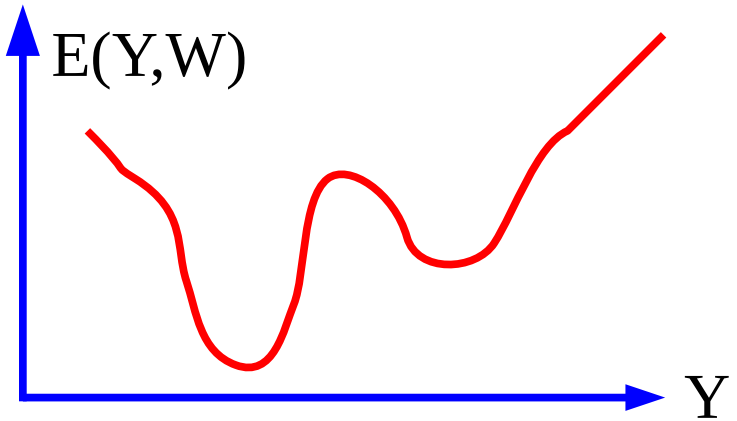

Learning the Energy Function

- parameterized energy function E(Y,W)
  - Make the energy low on the samples
  - Make the energy higher everywhere else
  - Making the energy low on the samples is easy
  - But how do we make it higher everywhere else?

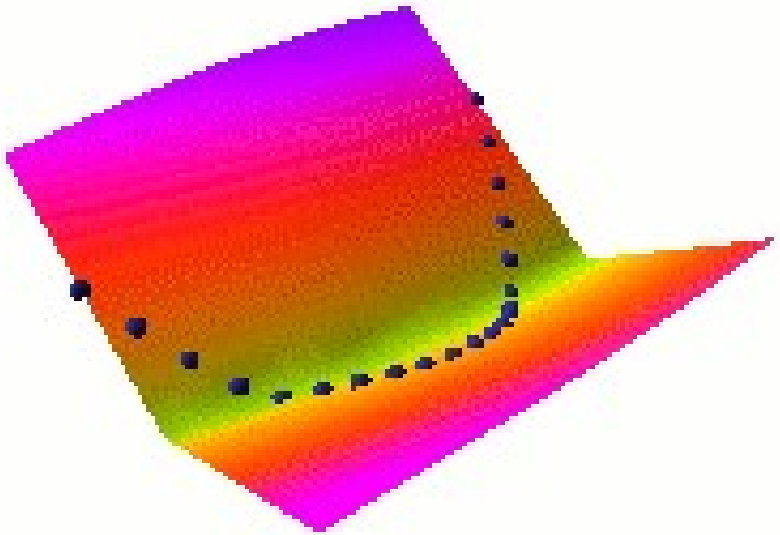

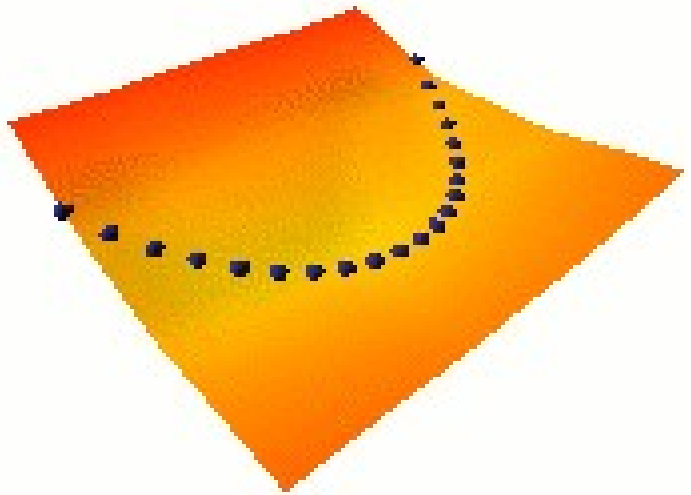

Seven Strategies to Shape the Energy Function

- 1. build the machine so that the volume of low energy stuff is constant
  - PCA, K-means, GMM, square ICA
- 2. push down of the energy of data points, push up everywhere else
  - Max likelihood (needs tractable partition function)
- 3. push down of the energy of data points, push up on chosen locations
  - contrastive divergence, Ratio Matching, Noise Contrastive Estimation, Minimum Probability Flow
- 4. minimize the gradient and maximize the curvature around data points
  - score matching
- 5. train a dynamical system so that the dynamics goes to the manifold
  - denoising auto-encoder
- 6. use a regularizer that limits the volume of space that has low energy
  - Sparse coding, sparse auto-encoder, PSD
- 7. if E(Y) = ||Y - G(Y)||^2, make G(Y) as "constant" as possible.
  - Contracting auto-encoder, saturating auto-encoder

#1: constant volume of low energy
Energy surface for PCA and K-means

- 1. build the machine so that the volume of low energy stuff is constant
  - PCA, K-means, GMM, square ICA...

PCA

$$E(Y) = \|W^T WY - Y\|^2$$

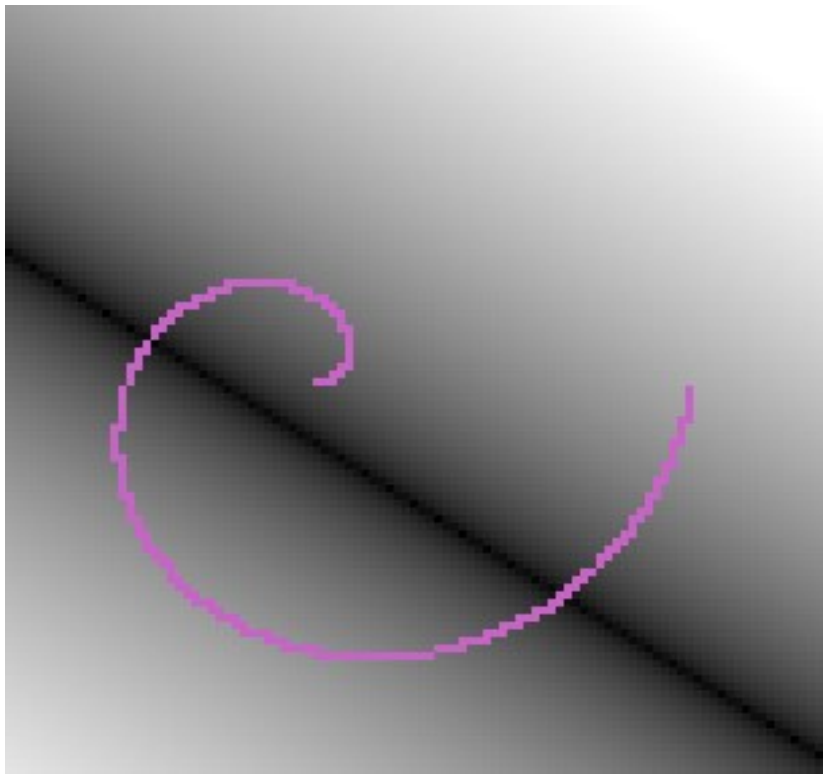

K-Means,
Z constrained to 1-of-K code

$$E(Y) = min_z \sum_i \|Y - W_i Z_i\|^2$$

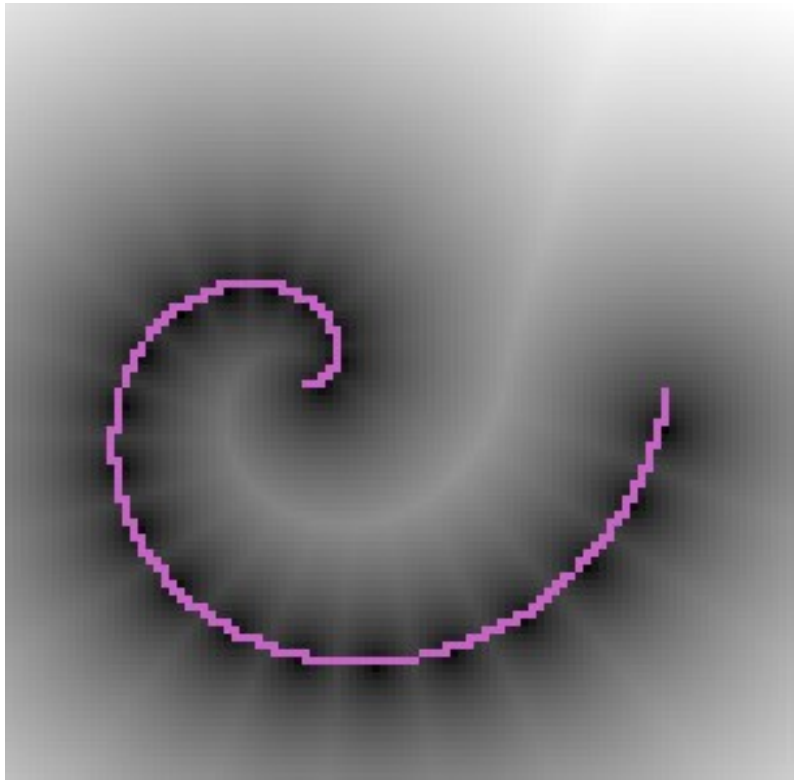

#2: push down of the energy of data points, push up everywhere else

- Max likelihood (requires a tractable partition function)

Maximizing P(Y|W) on training samples

make this big

$$P(Y|W) = \frac{e^{-\beta E(Y,W)}}{\int_y e^{-\beta E(y,W)}}$$

make this small

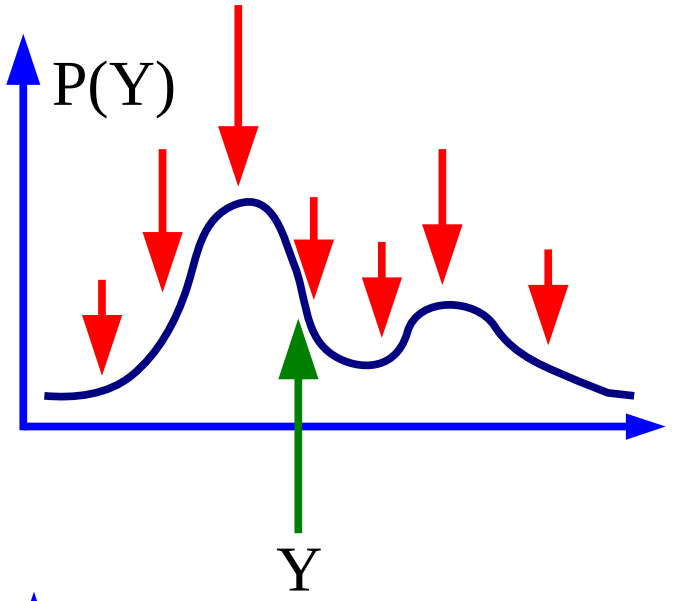

Minimizing -log P(Y,W) on training samples

$$L(Y,W) = E(Y,W) + \frac{1}{\beta} \log \int_y e^{-\beta E(y,W)}$$

make this small

make this big

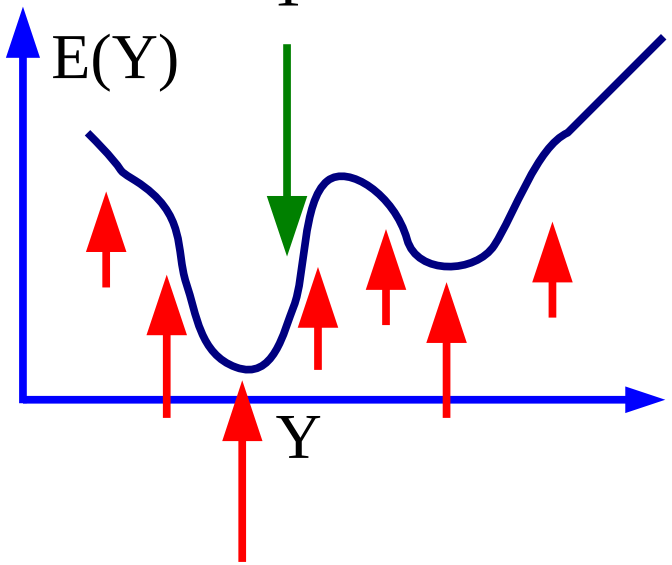

#2: push down of the energy of data points, push up everywhere else

Gradient of the negative log-likelihood loss for one sample Y:

$$\frac{\partial L(Y,W)}{\partial W} = \frac{\partial E(Y,W)}{\partial W} - \int_y P(y|W) \frac{\partial E(y,W)}{\partial W}$$

Gradient descent:

$$W \leftarrow W - \eta \frac{\partial L(Y,W)}{\partial W}$$

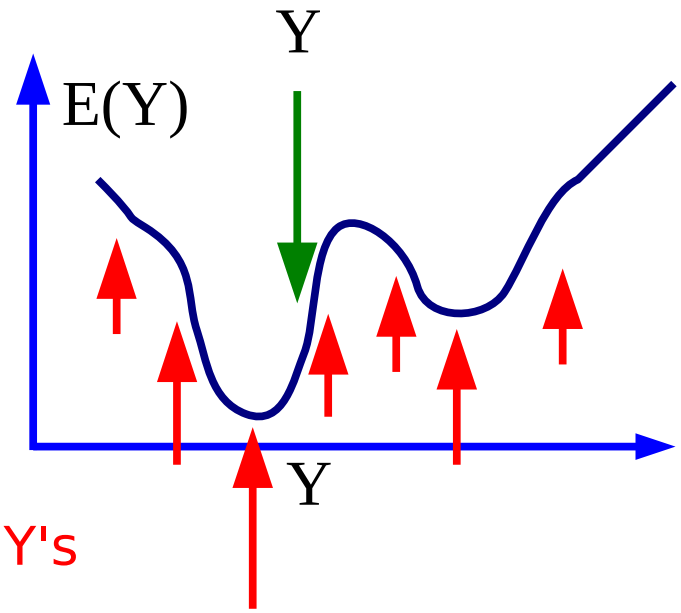

Pushes down on the energy of the samples

Pulls up on the energy of low-energy Y's

$$W \leftarrow W - \eta \frac{\partial E(Y,W)}{\partial W} + \eta \int_y P(y|W) \frac{\partial E(y,W)}{\partial W}$$

#3. push down of the energy of data points, push up on chosen locations

- contrastive divergence, Ratio Matching, Noise Contrastive Estimation, Minimum Probability Flow
- Contrastive divergence: basic idea
  - Pick a training sample, lower the energy at that point
  - From the sample, move down in the energy surface with noise
  - Stop after a while
  - Push up on the energy of the point where we stopped
  - This creates grooves in the energy surface around data manifolds
  - CD can be applied to any energy function (not just RBMs)
- Persistent CD: use a bunch of "particles" and remember their positions
  - Make them roll down the energy surface with noise
  - Push up on the energy wherever they are
  - Faster than CD
- RBM

$$E(Y,Z) = -Z^T WY \qquad E(Y) = -\log \sum_z e^{Z^T WY}$$

#6. use a regularizer that limits the volume of space that has low energy

- Sparse coding, sparse auto-encoder, Predictive Sparse Decomposition

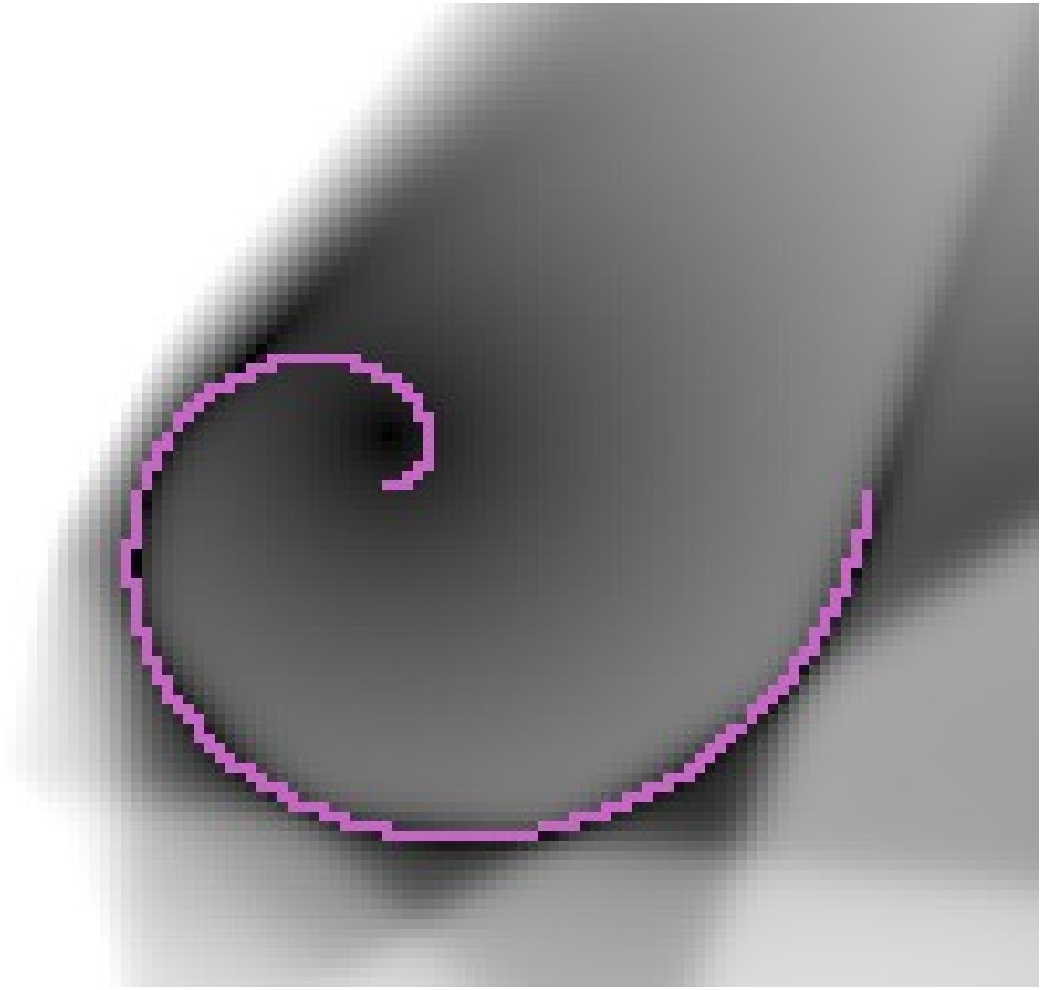

Energy Functions of Various Methods

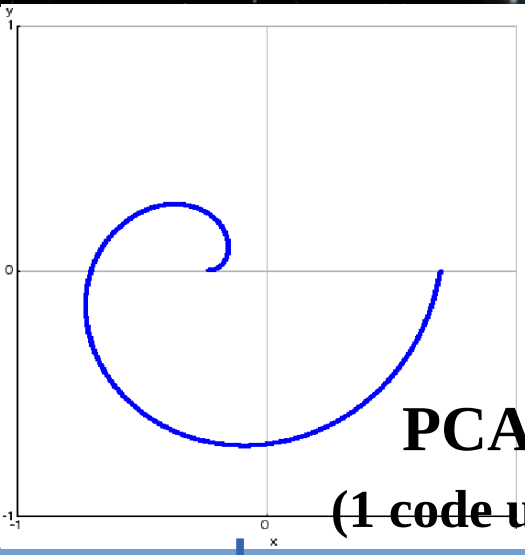

- 2 dimensional toy dataset: spiral
- Visualizing energy surface
  - (black = low, white = high)

	PCA (1 code unit)	autoencoder (1 code unit)	sparse coding (20 code units)	K-Means (20 code units)
encoder	$W'Y$	$\sigma(W_e Y)$	$\sigma(W_e Z)$	–
decoder	WZ	$W_d Z$	$W_d Z$	WZ
energy	$\|\|Y-WZ\|\|^2$	$\|\|Y-WZ\|\|^2$	$\|\|Y-WZ\|\|^2$	$\|\|Y-WZ\|\|^2$
loss	$F(Y)$	$F(Y)$	$F(Y)$	$F(Y)$
pull-up	dimens.	dimens.	sparsity	1-of-N code
	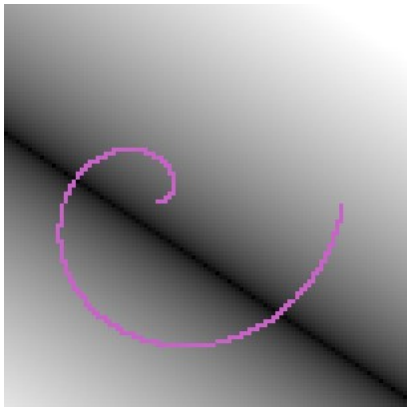	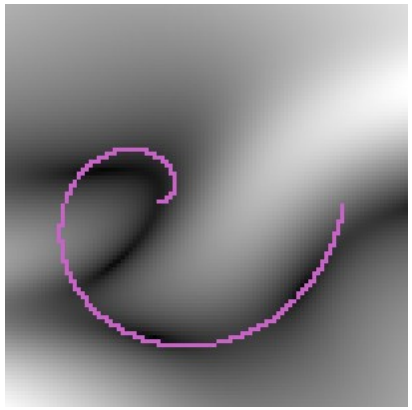	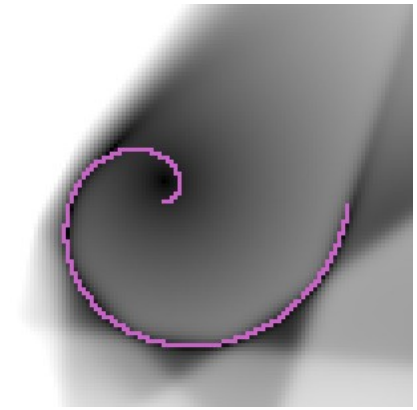	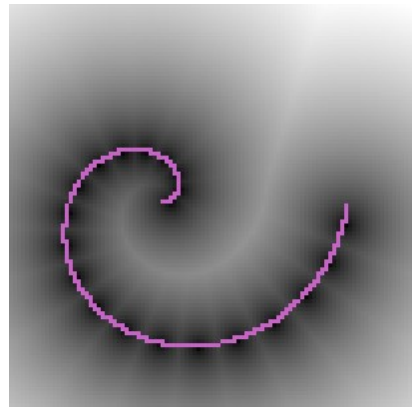

Dictionary Learning With
Fast Approximate Inference:
Sparse Auto-Encoders

How to Speed Up Inference in a Generative Model?

- Factor Graph with an asymmetric factor
- Inference Z → Y is easy
  - Run Z through deterministic decoder, and sample Y
- Inference Y → Z is hard, particularly if Decoder function is many-to-one
  - MAP: minimize sum of two factors with respect to Z
  - Z* = argmin_z Distance[Decoder(Z), Y] + FactorB(Z)
- Examples: K-Means (1of K), Sparse Coding (sparse), Factor Analysis

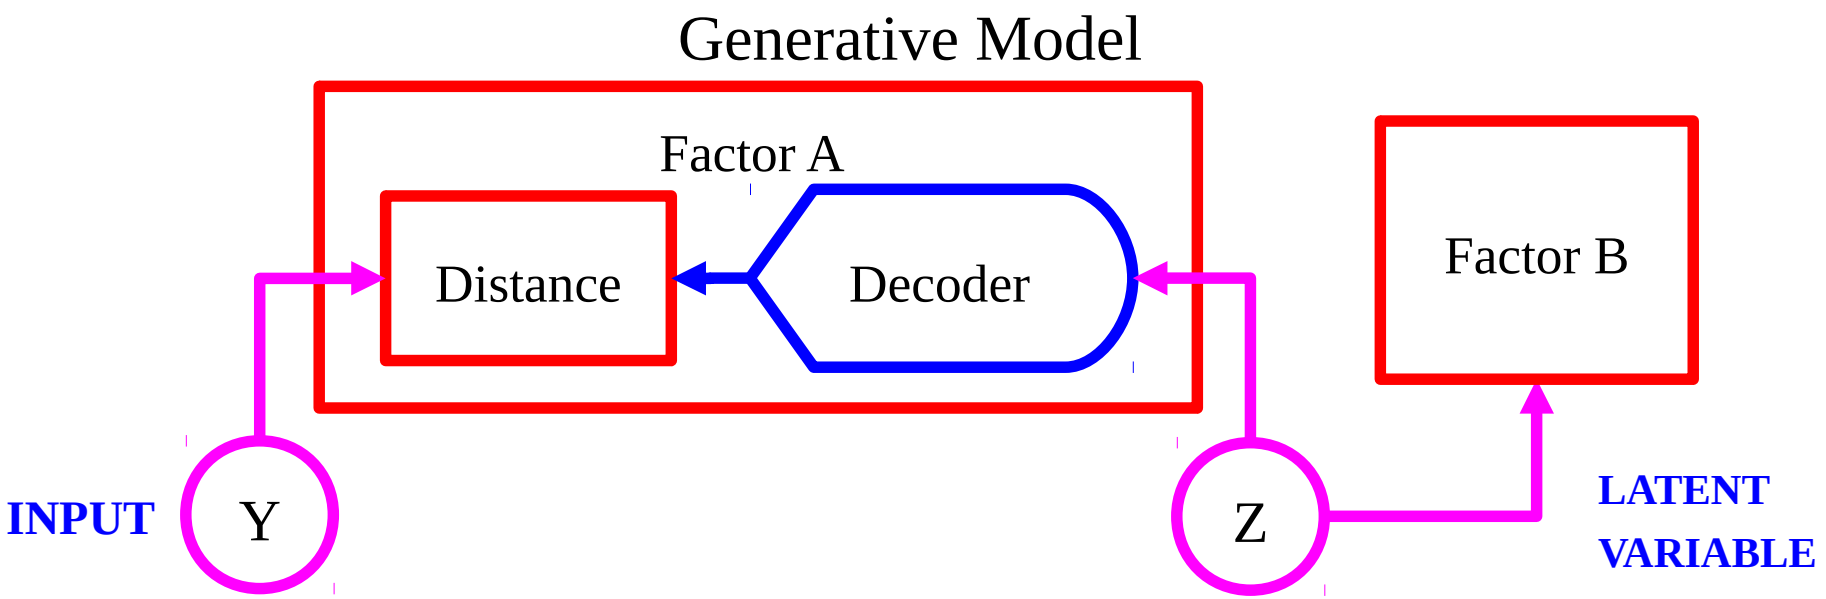

Sparse Modeling: Sparse Coding + Dictionary Learning

[Olshausen & Field 1997]

- Sparse linear reconstruction
- Energy = reconstruction_error + code_prediction_error + code_sparsity

$$E(Y^i, Z) = \|Y^i - W_d Z\|^2 + \lambda \sum_j |z_j|$$

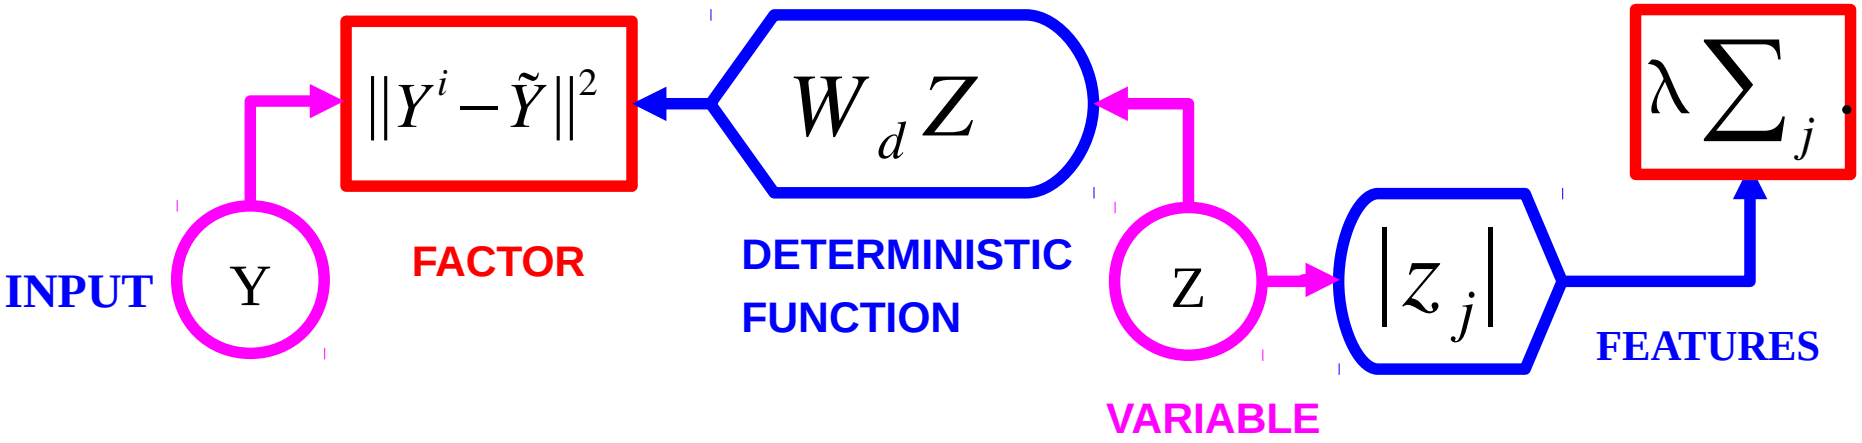

- Inference is expensive: ISTA/FISTA, CGIHT, coordinate descent....

$$Y \rightarrow \hat{Z} = argmin_Z E(Y, Z)$$

#6. use a regularizer that limits the volume of space that has low energy

- Sparse coding, sparse auto-encoder, Predictive Saprse Decomposition

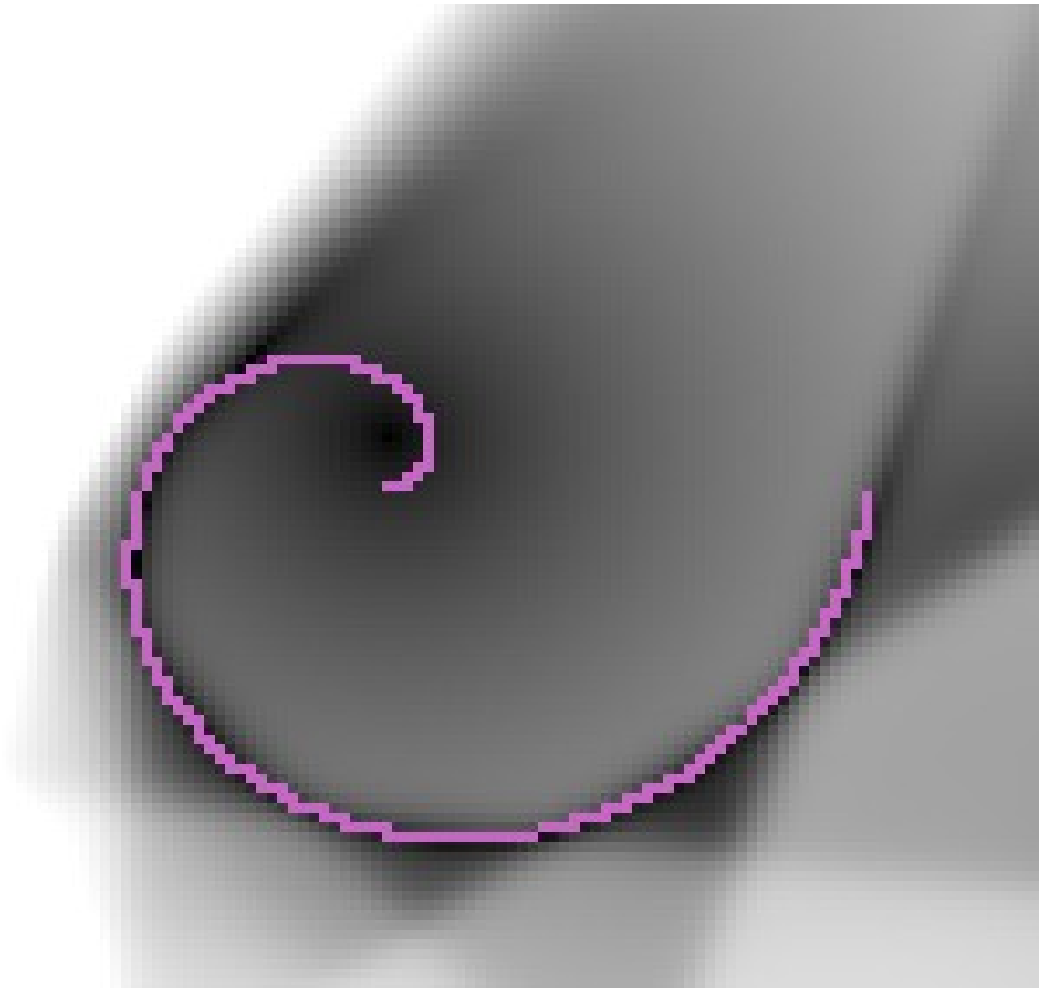

Encoder-Decoder Architecture

[Kavukcuoglu, Ranzato, LeCun, rejected by every conference, 2008-2009]

- Train a "simple" feed-forward function to predict the result of a complex optimization on the data points of interest

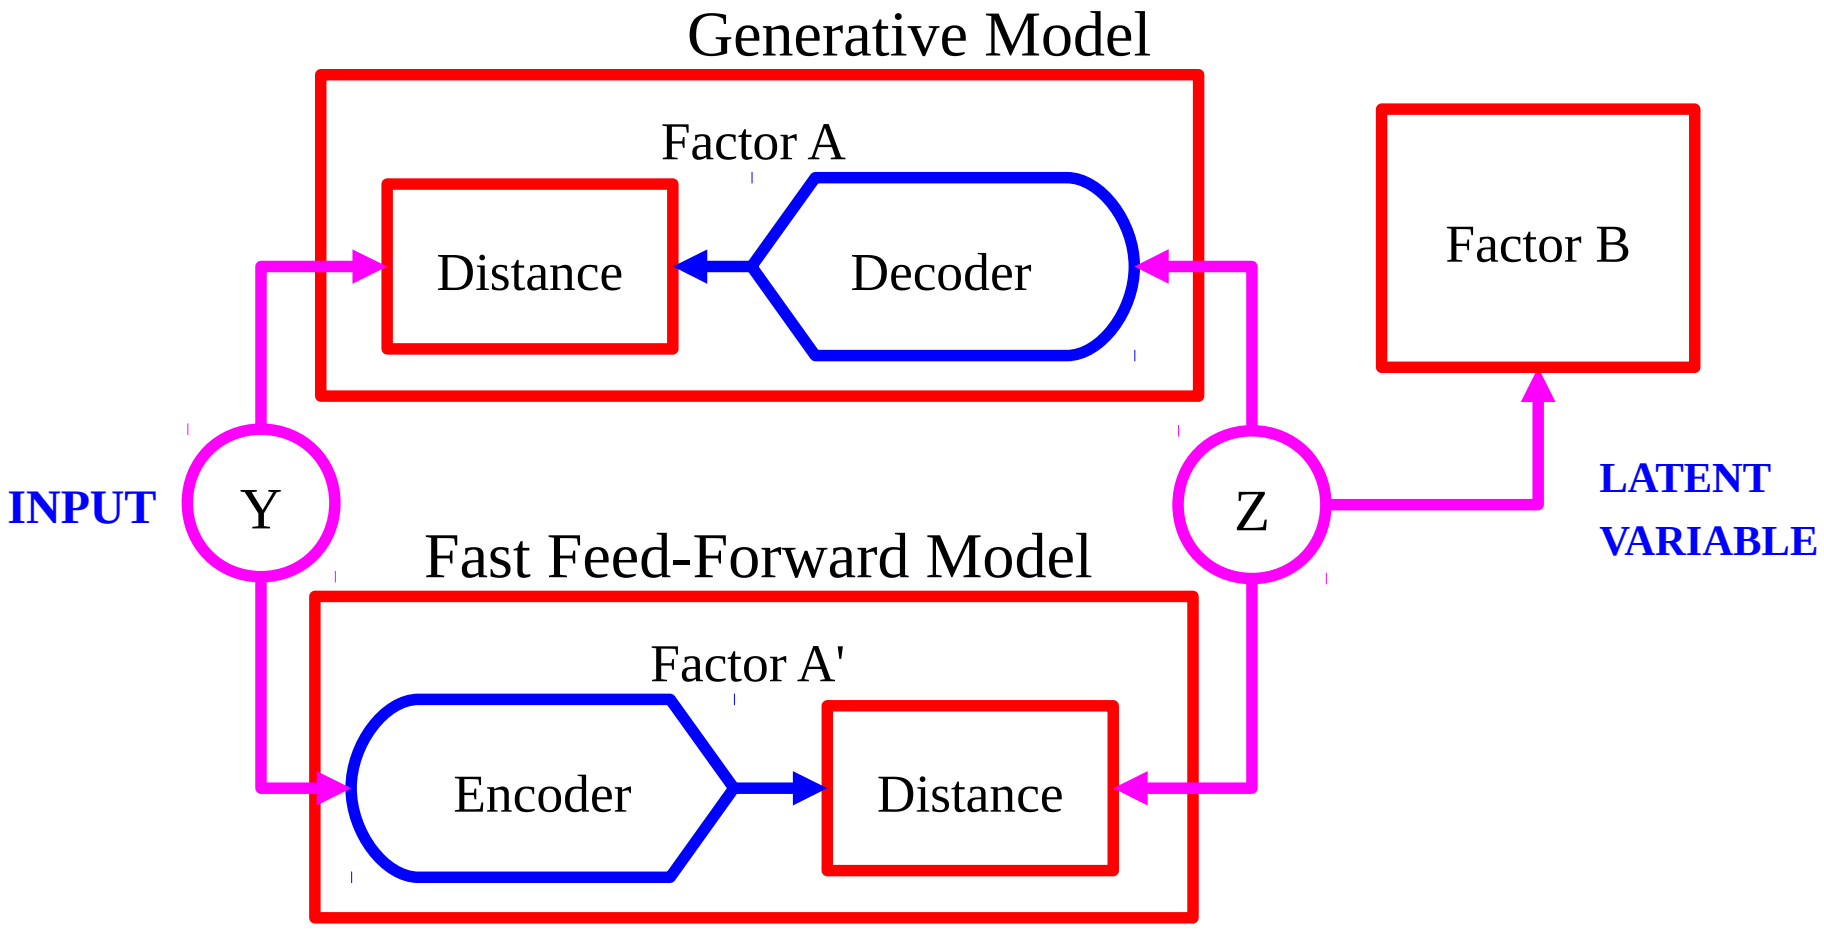

- 1. Find optimal Zi for all Yi; 2. Train Encoder to predict Zi from Yi

Learning to Perform
Approximate Inference:
Predictive Sparse Decomposition
Sparse Auto-Encoders

Sparse auto-encoder: Predictive Sparse Decomposition (PSD)

[Kavukcuoglu, Ranzato, LeCun, 2008 → arXiv:1010.3467],

- Prediction the optimal code with a trained encoder
- Energy = reconstruction_error + code_prediction_error + code_sparsity

$$E(Y^i, Z) = \|Y^i - W_d Z\|^2 + \|Z - g_e(W_e, Y^i)\|^2 + \lambda \sum_j |z_j|$$

$$g_e(W_e, Y^i) = shrinkage(W_e Y^i)$$

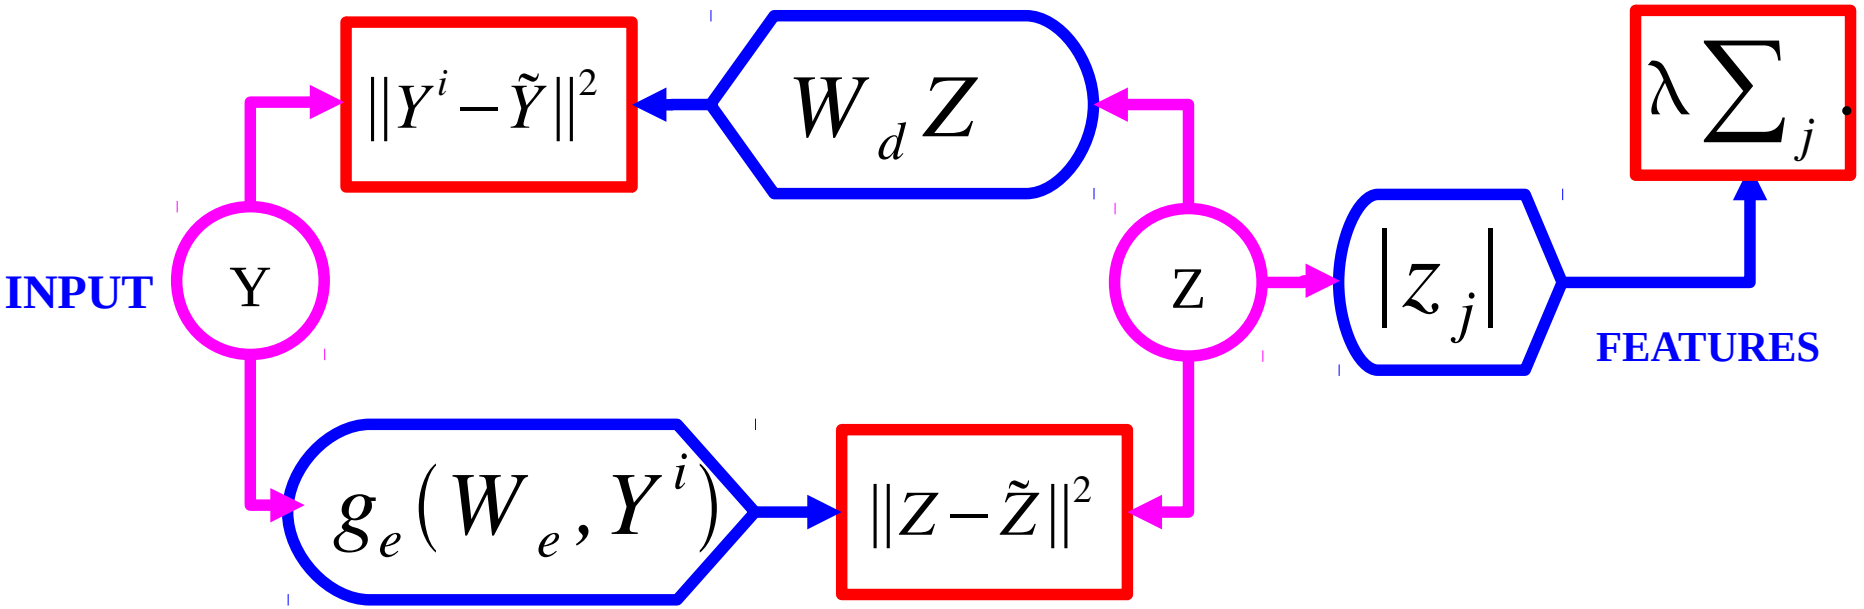

Regularized Encoder-Decoder Model (auto-Encoder) for Unsupervised Feature Learning

- **Encoder:** computes feature vector Z from input X
- **Decoder:** reconstructs input X from feature vector Z
- **Feature vector:** high dimensional and regularized (e.g. sparse)
- **Factor graph with energy function E(X,Z) with 3 terms:**
  - Linear decoding function and reconstruction error
  - Non-Linear encoding function and prediction error term
  - Pooling function and regularization term (e.g. sparsity)

$$E(Y,Z) = \|Y - W_d Z\|^2 + \|Z - g_e(W_e, Y)\|^2 + \sum_j \sqrt{\sum_{k \in P_j} Z_k^2}$$

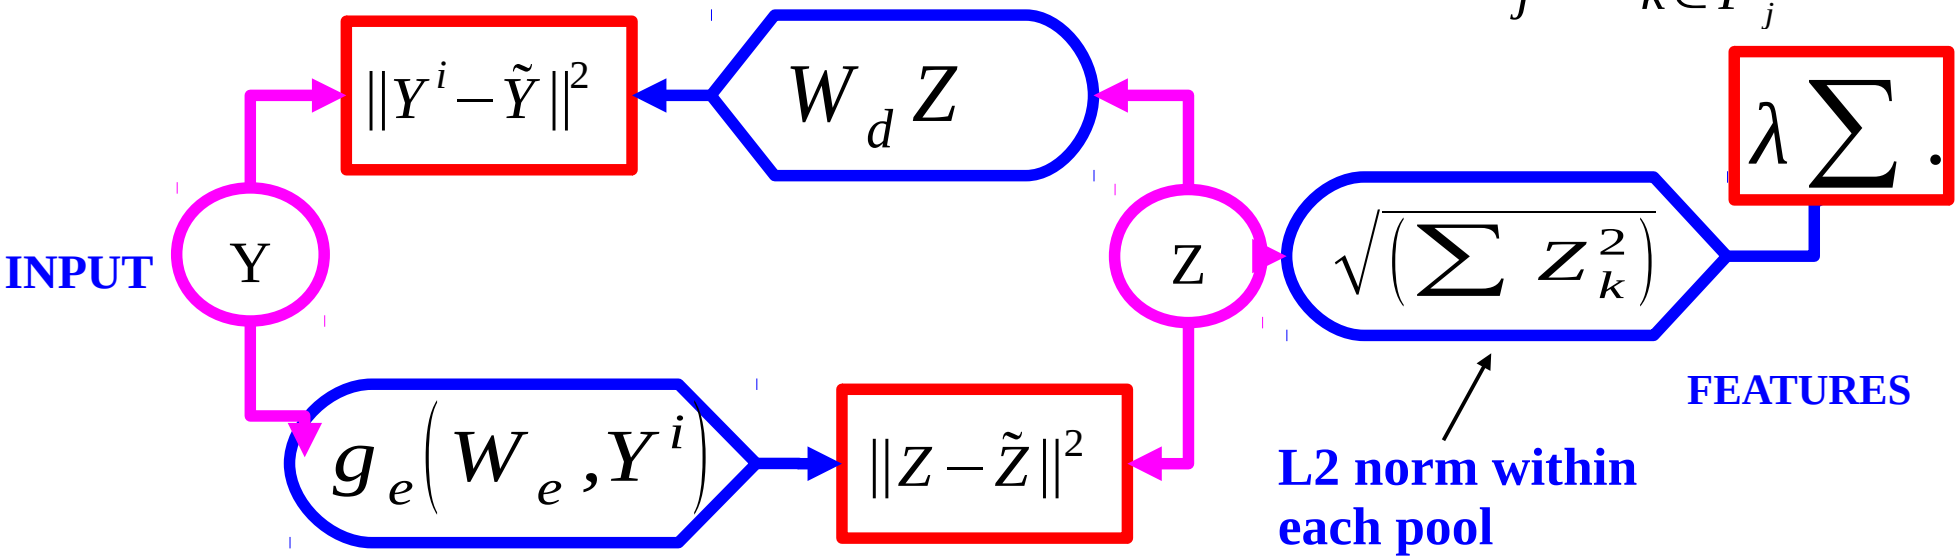

PSD: Basis Functions on MNIST

- Basis functions (and encoder matrix) are digit parts

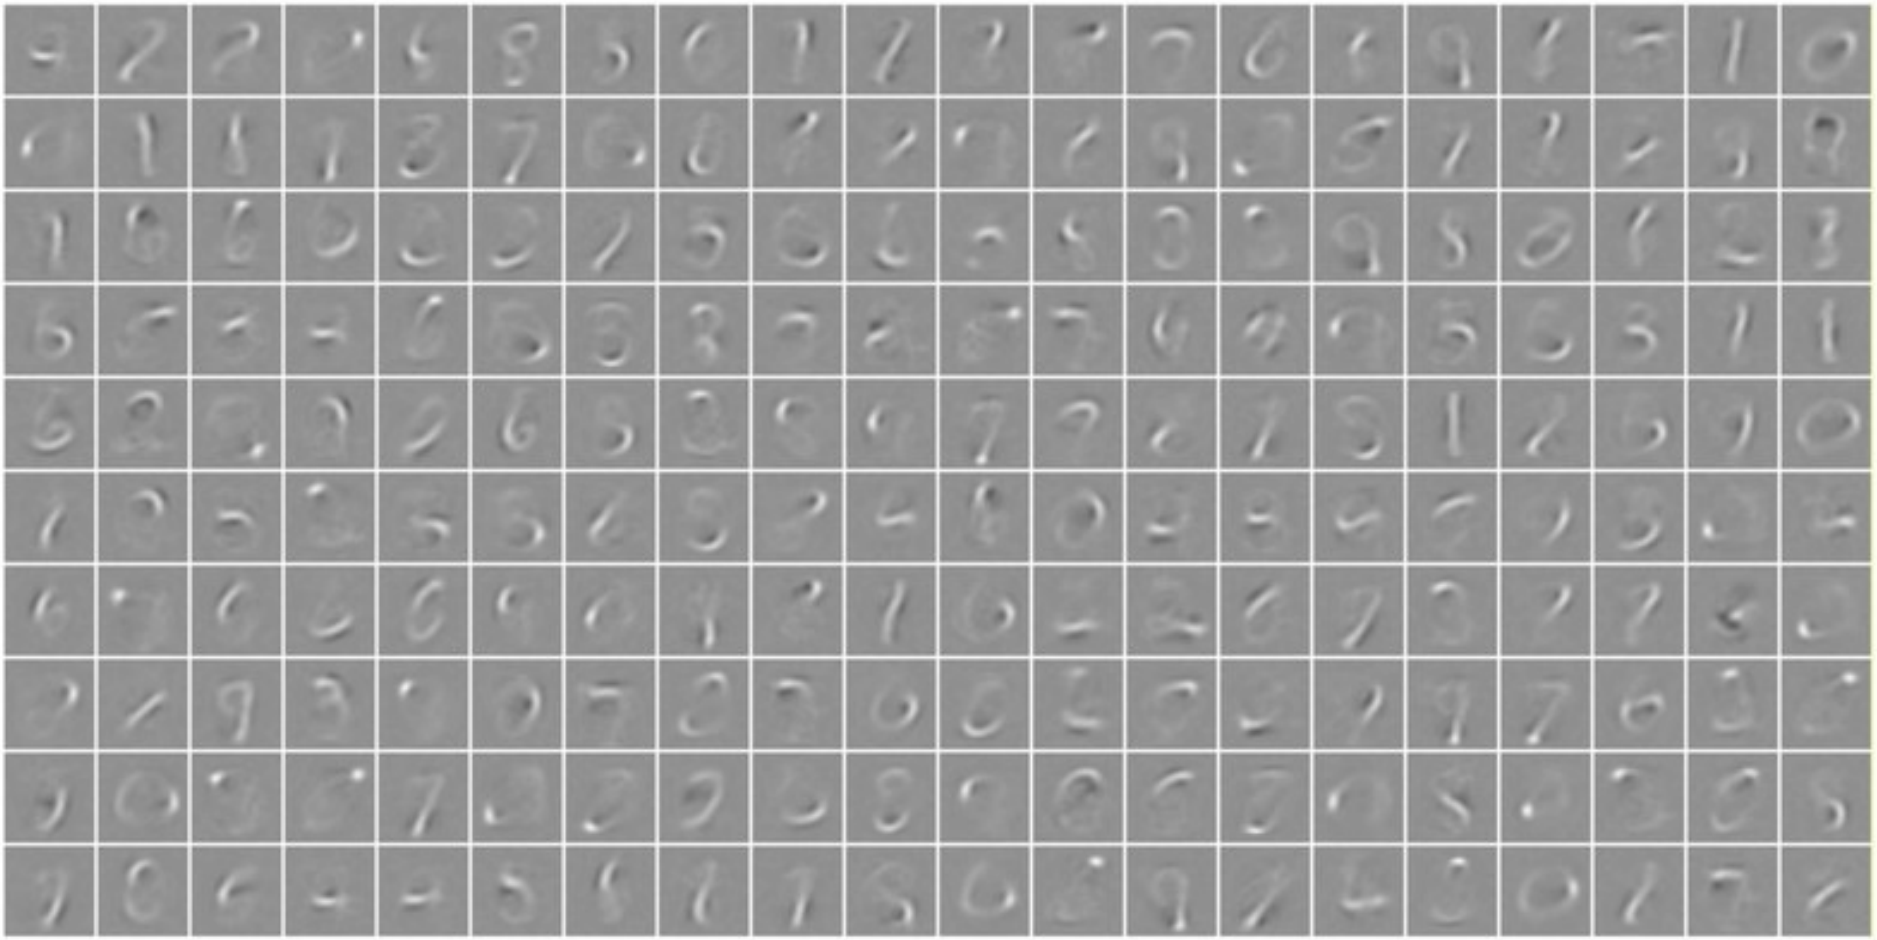

Predictive Sparse Decomposition (PSD): Training

- Training on natural images patches.
  - 12X12
  - 256 basis functions

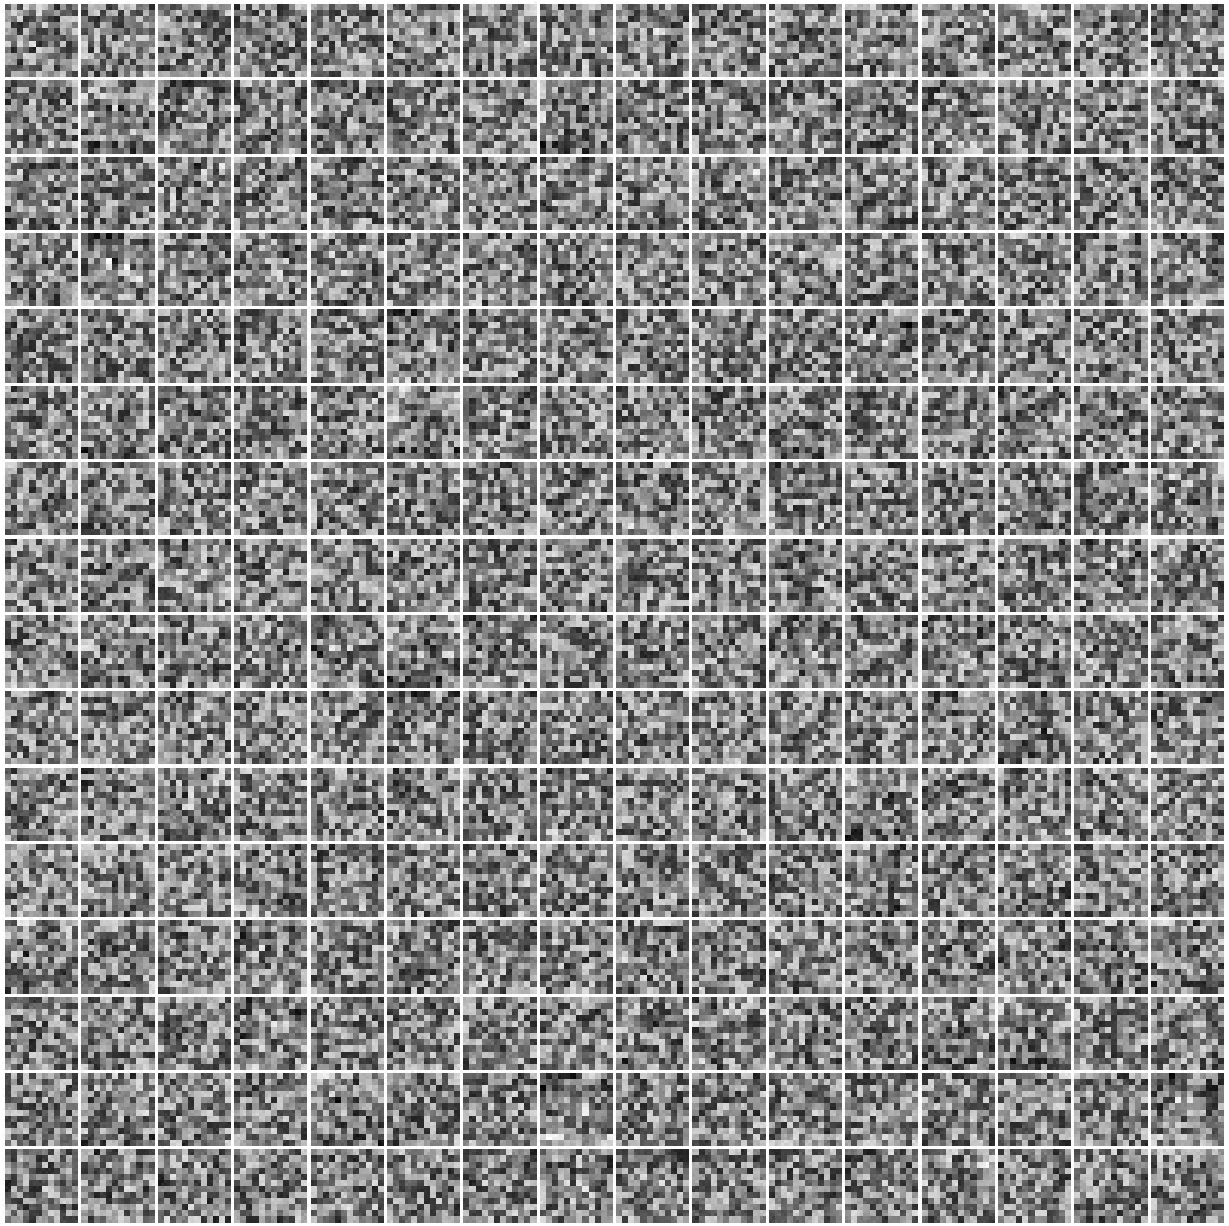

Learned Features on natural patches: V1-like receptive fields

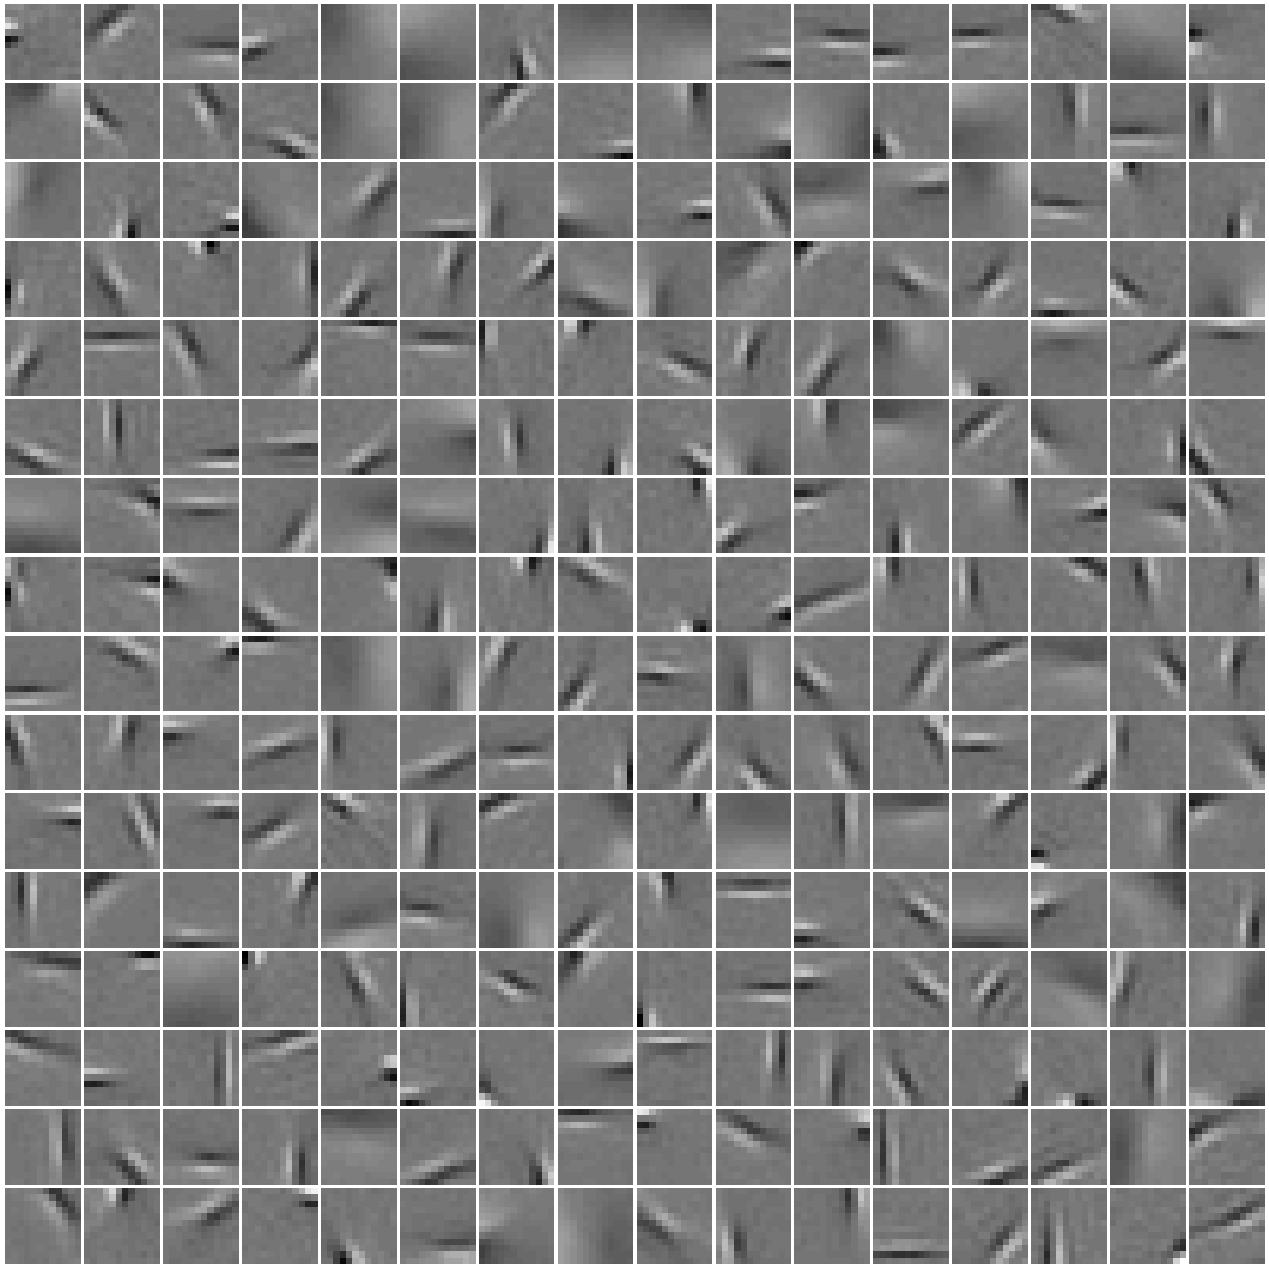

Learning to Perform
Approximate Inference
LISTA

Better Idea: Give the "right" structure to the encoder

- ISTA/FISTA: iterative algorithm that converges to optimal sparse code

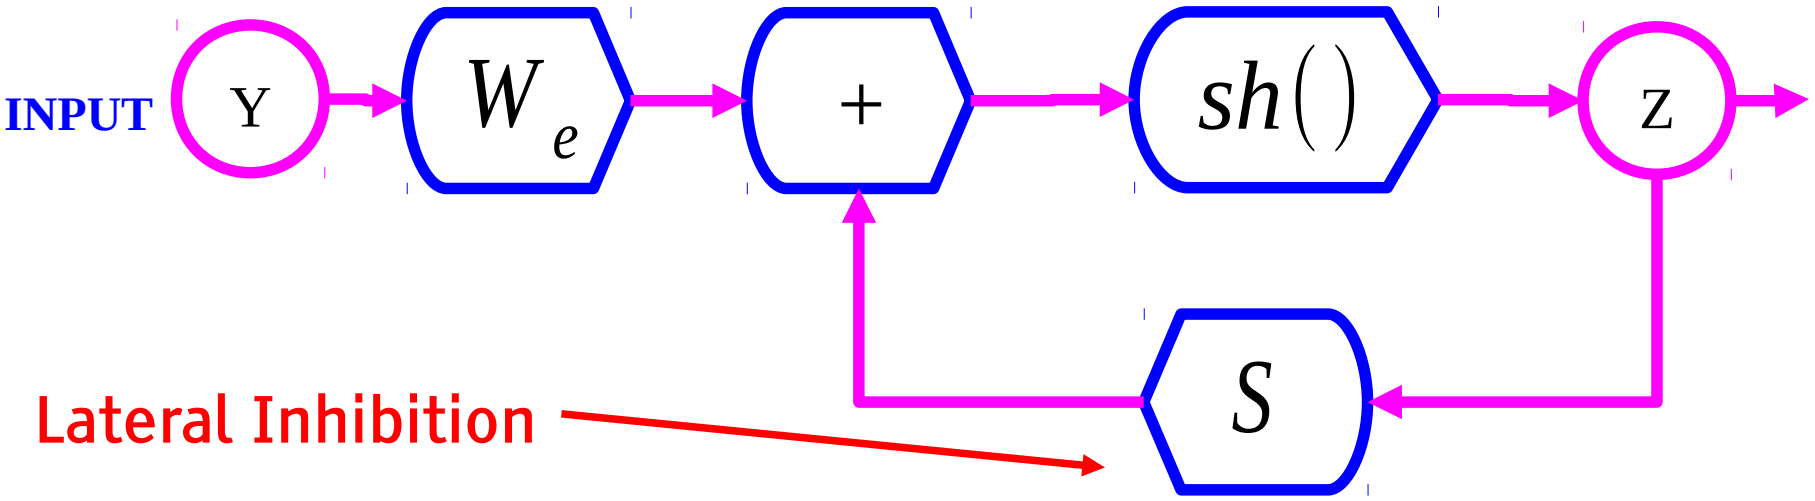

$$Z(t+1) = \text{Shrinkage}_{\lambda/L}\left[Z(t) - \frac{1}{L}W_d^T(W_d Z(t) - Y)\right]$$

- ISTA/FISTA reparameterized:

$$Z(t+1) = \text{Shrinkage}_{\lambda/L}\left[W_e^T Y + SZ(t)\right]; \quad W_e = \frac{1}{L}W_d; \quad S = I - \frac{1}{L}W_d^T W_d$$

- LISTA (Learned ISTA): learn the We and S matrices to get fast solutions

[Gregor & LeCun, ICML 2010], [Bronstein et al. ICML 2012], [Rolfe & LeCun ICLR 2013]

LISTA: Train We and S matrices to give a good approximation quickly

- Think of the FISTA flow graph as a recurrent neural net where We and S are trainable parameters

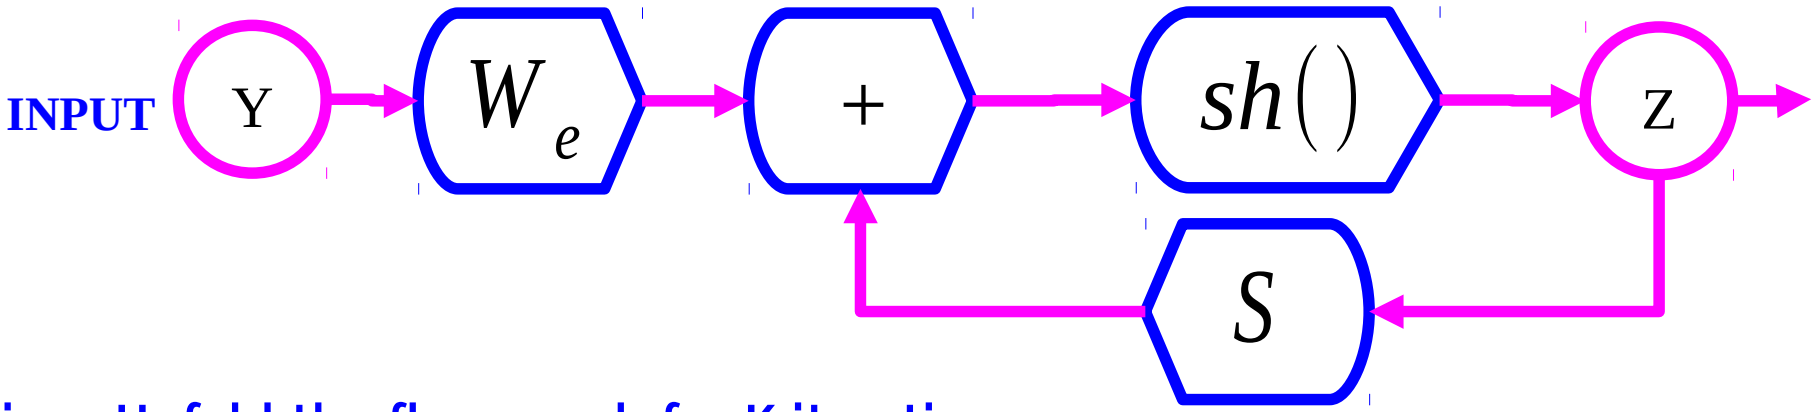

- Time-Unfold the flow graph for K iterations
- Learn the We and S matrices with "backprop-through-time"
- Get the best approximate solution within K iterations

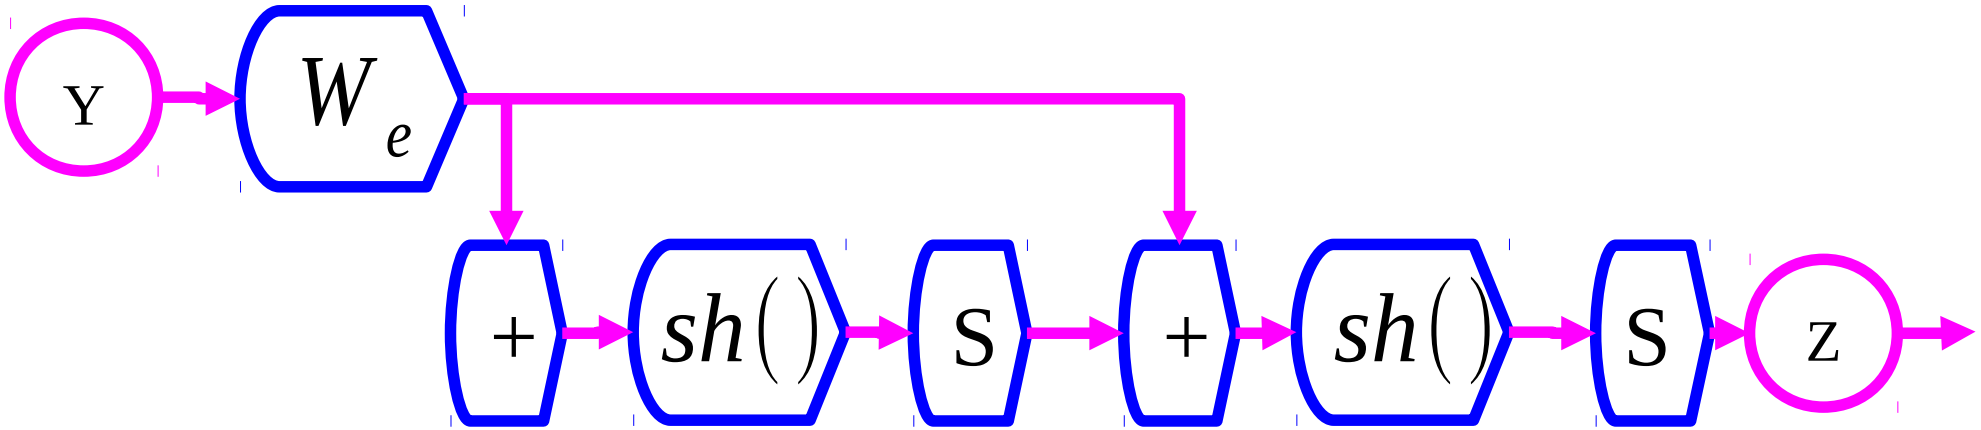

Learning ISTA (LISTA) vs ISTA/FISTA

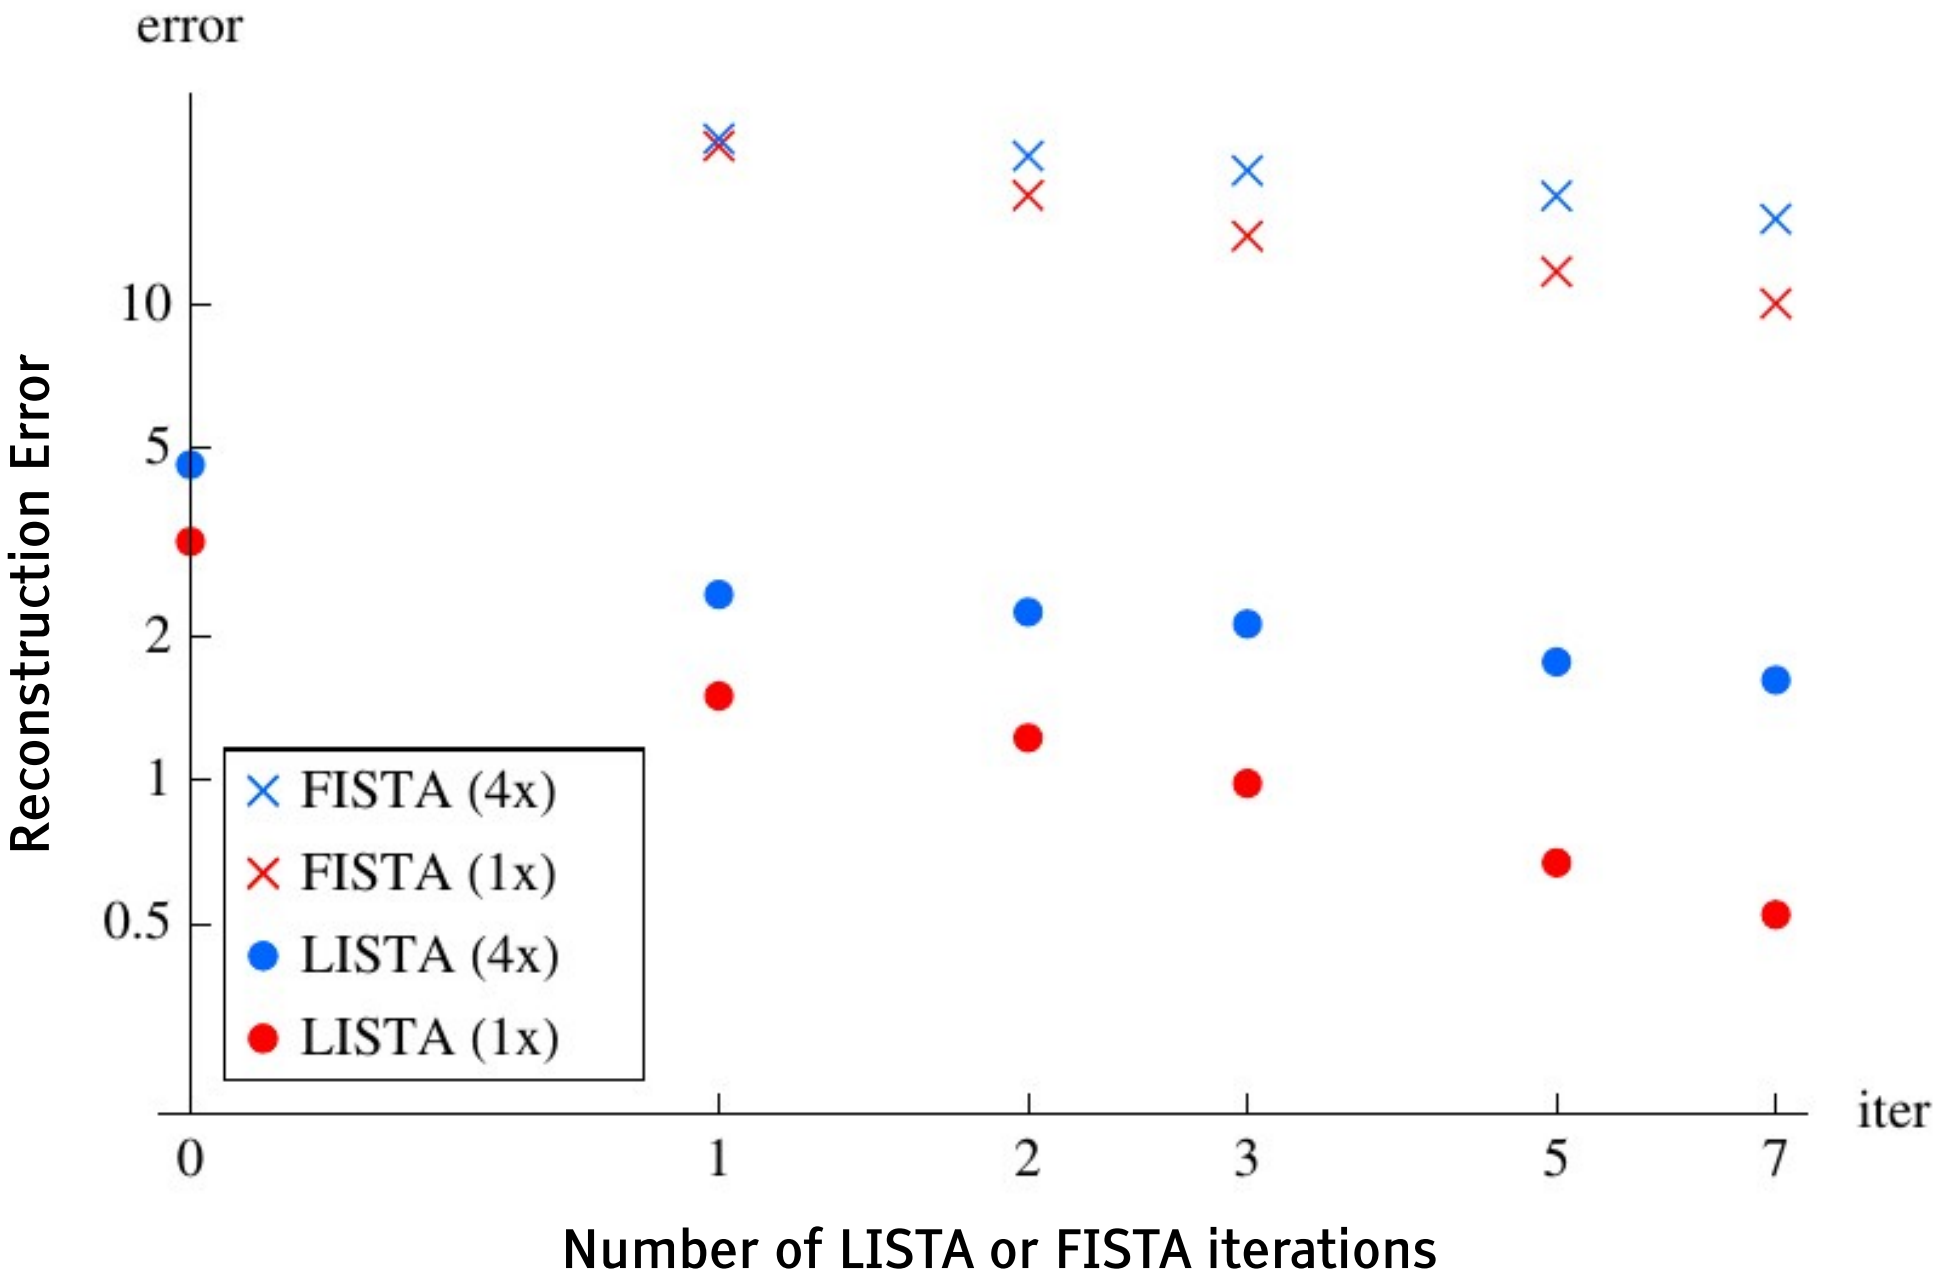

LISTA with partial mutual inhibition matrix

Y LeCun

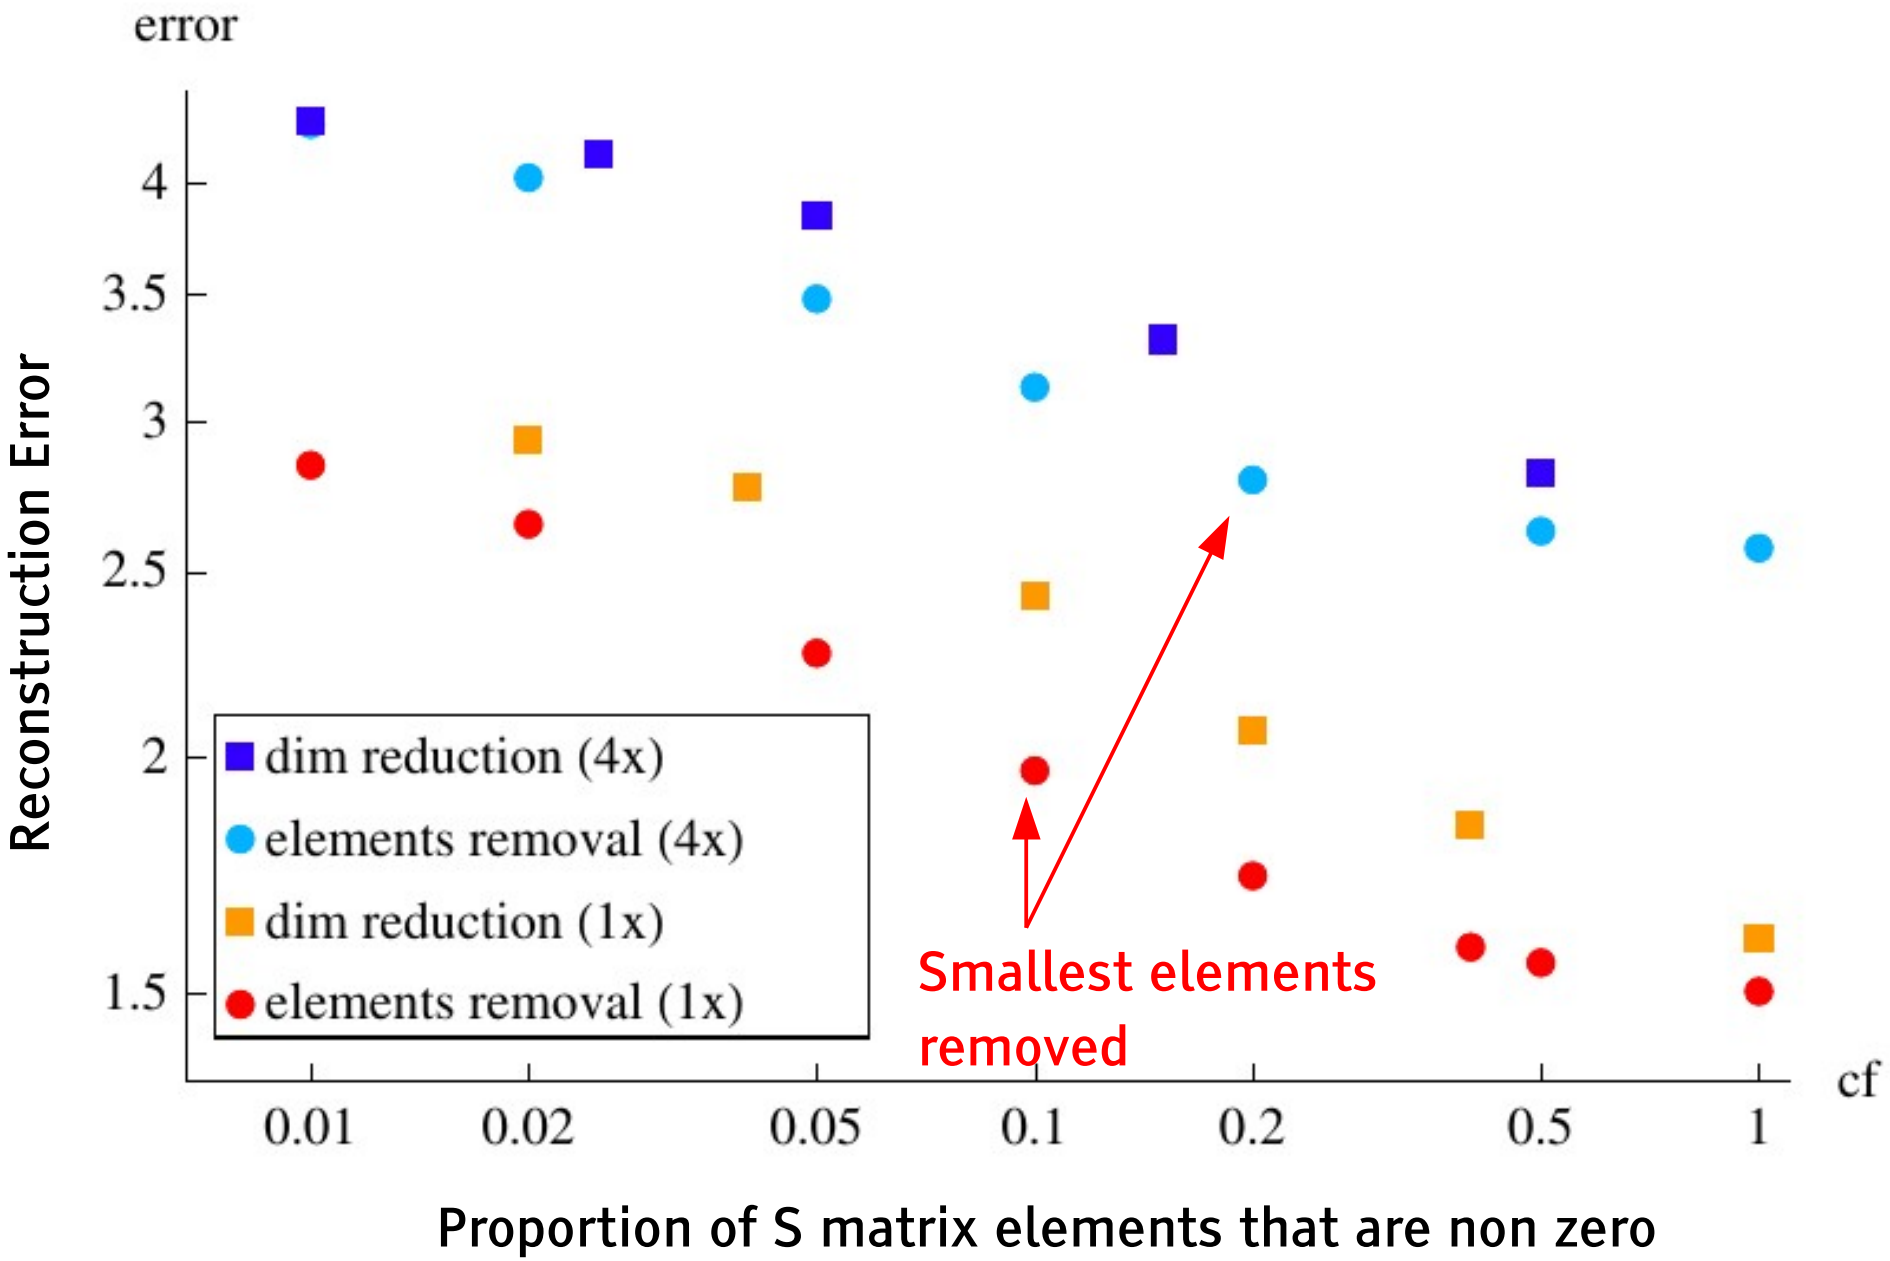

Learning Coordinate Descent (LcoD): faster than LISTA

Y LeCun

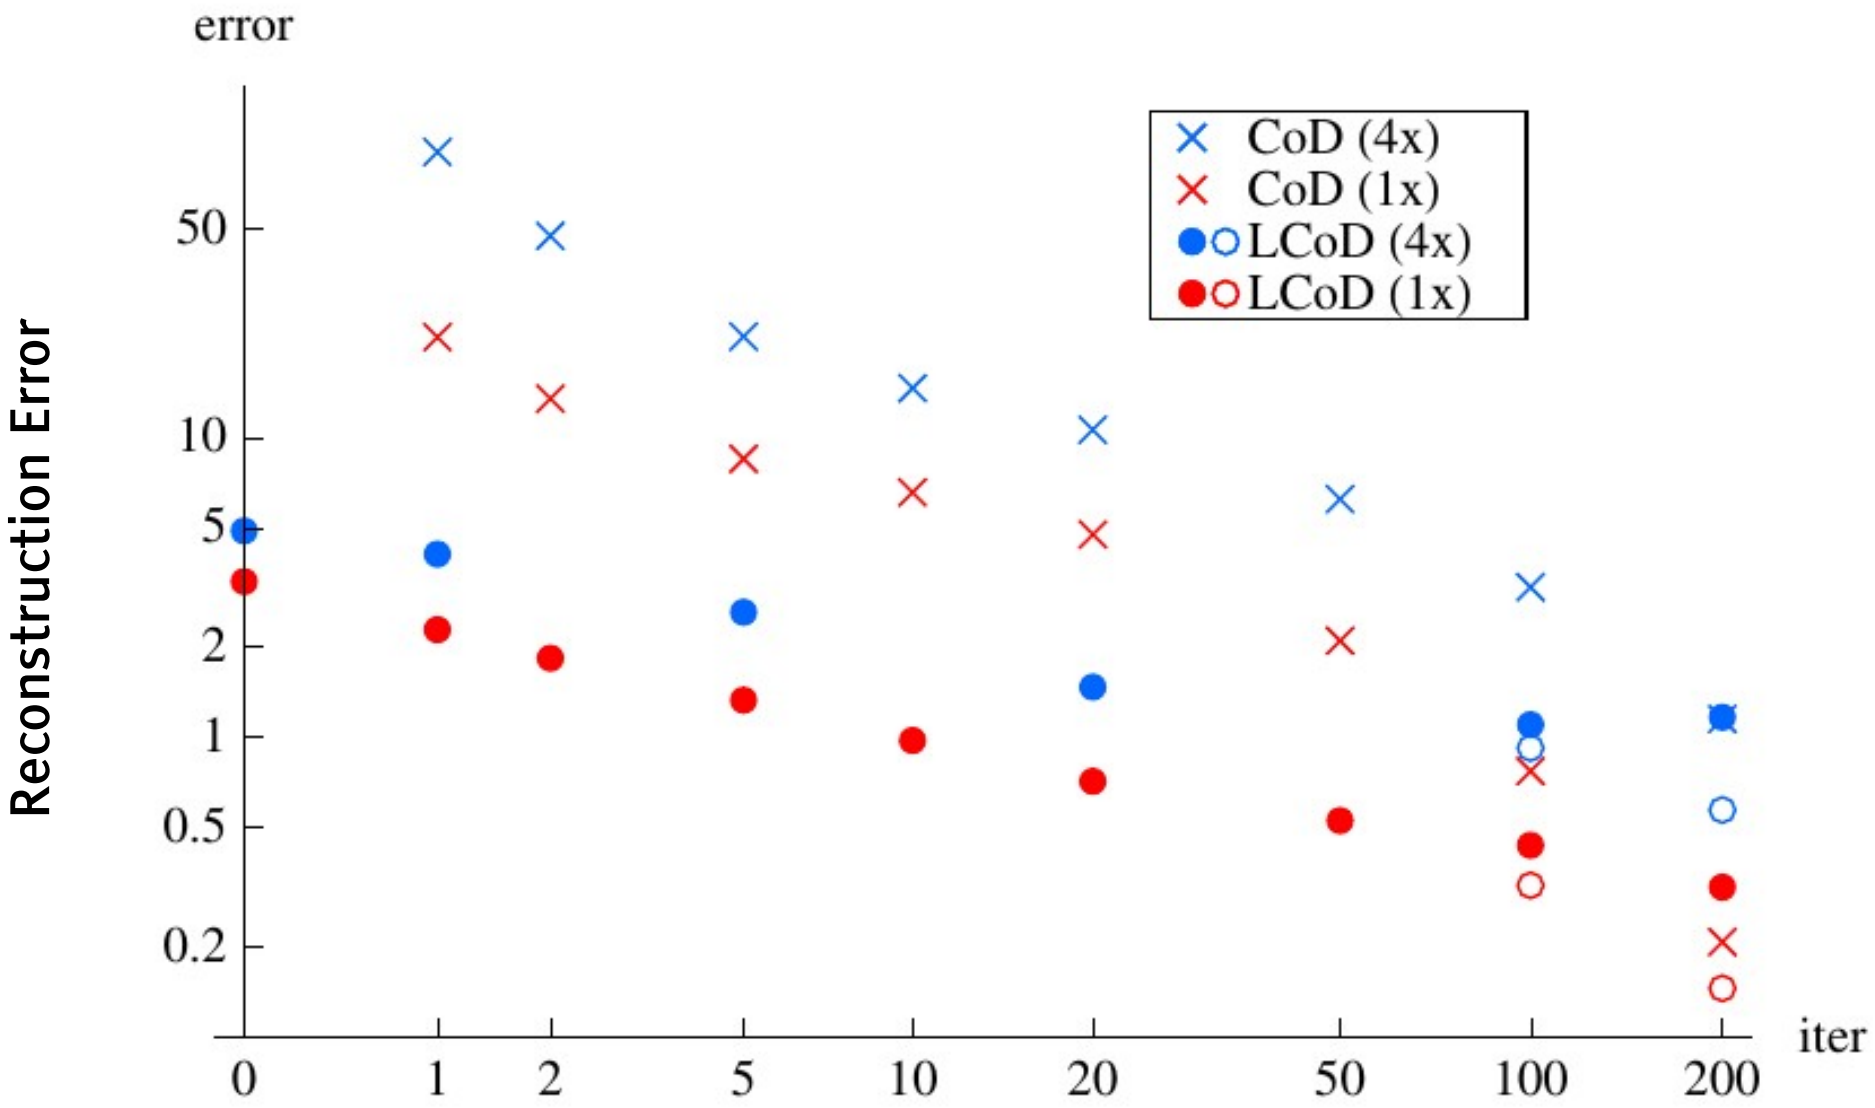

Discriminative Recurrent Sparse Auto-Encoder (DrSAE)

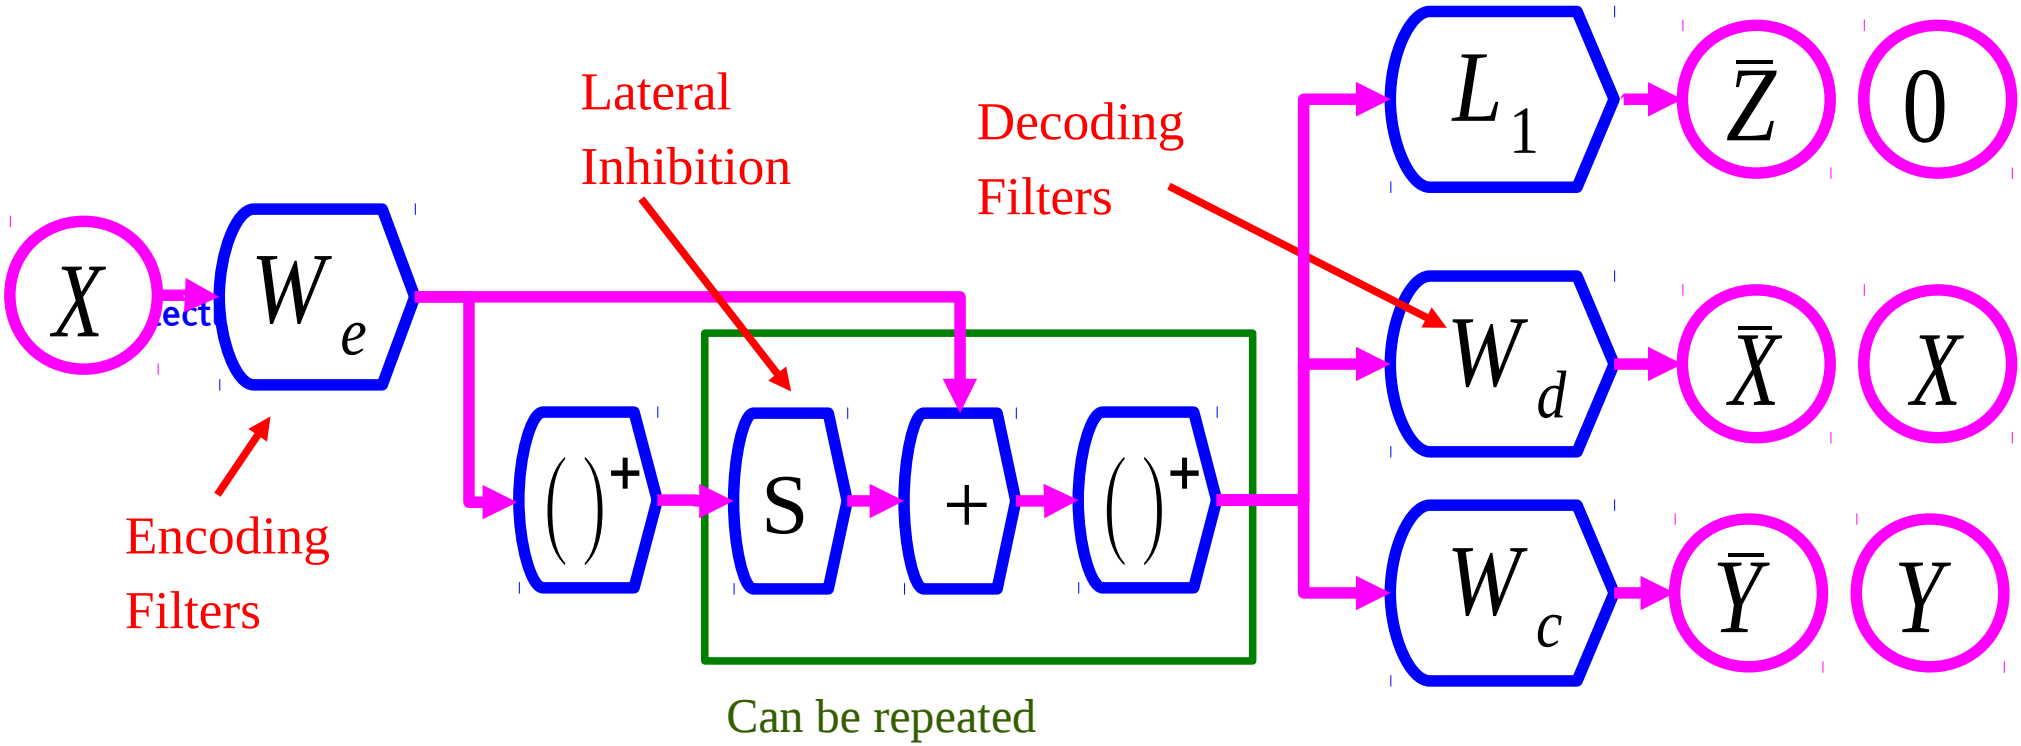

[Rolfe & LeCun ICLR 2013]

- Rectified linear units
- Classification loss: cross-entropy
- Reconstruction loss: squared error
- Sparsity penalty: L1 norm of last hidden layer
- Rows of Wd and columns of We constrained in unit sphere

DrSAE Discovers manifold structure of handwritten digits

Y LeCun

- Image = prototype + sparse sum of "parts" (to move around the manifold)

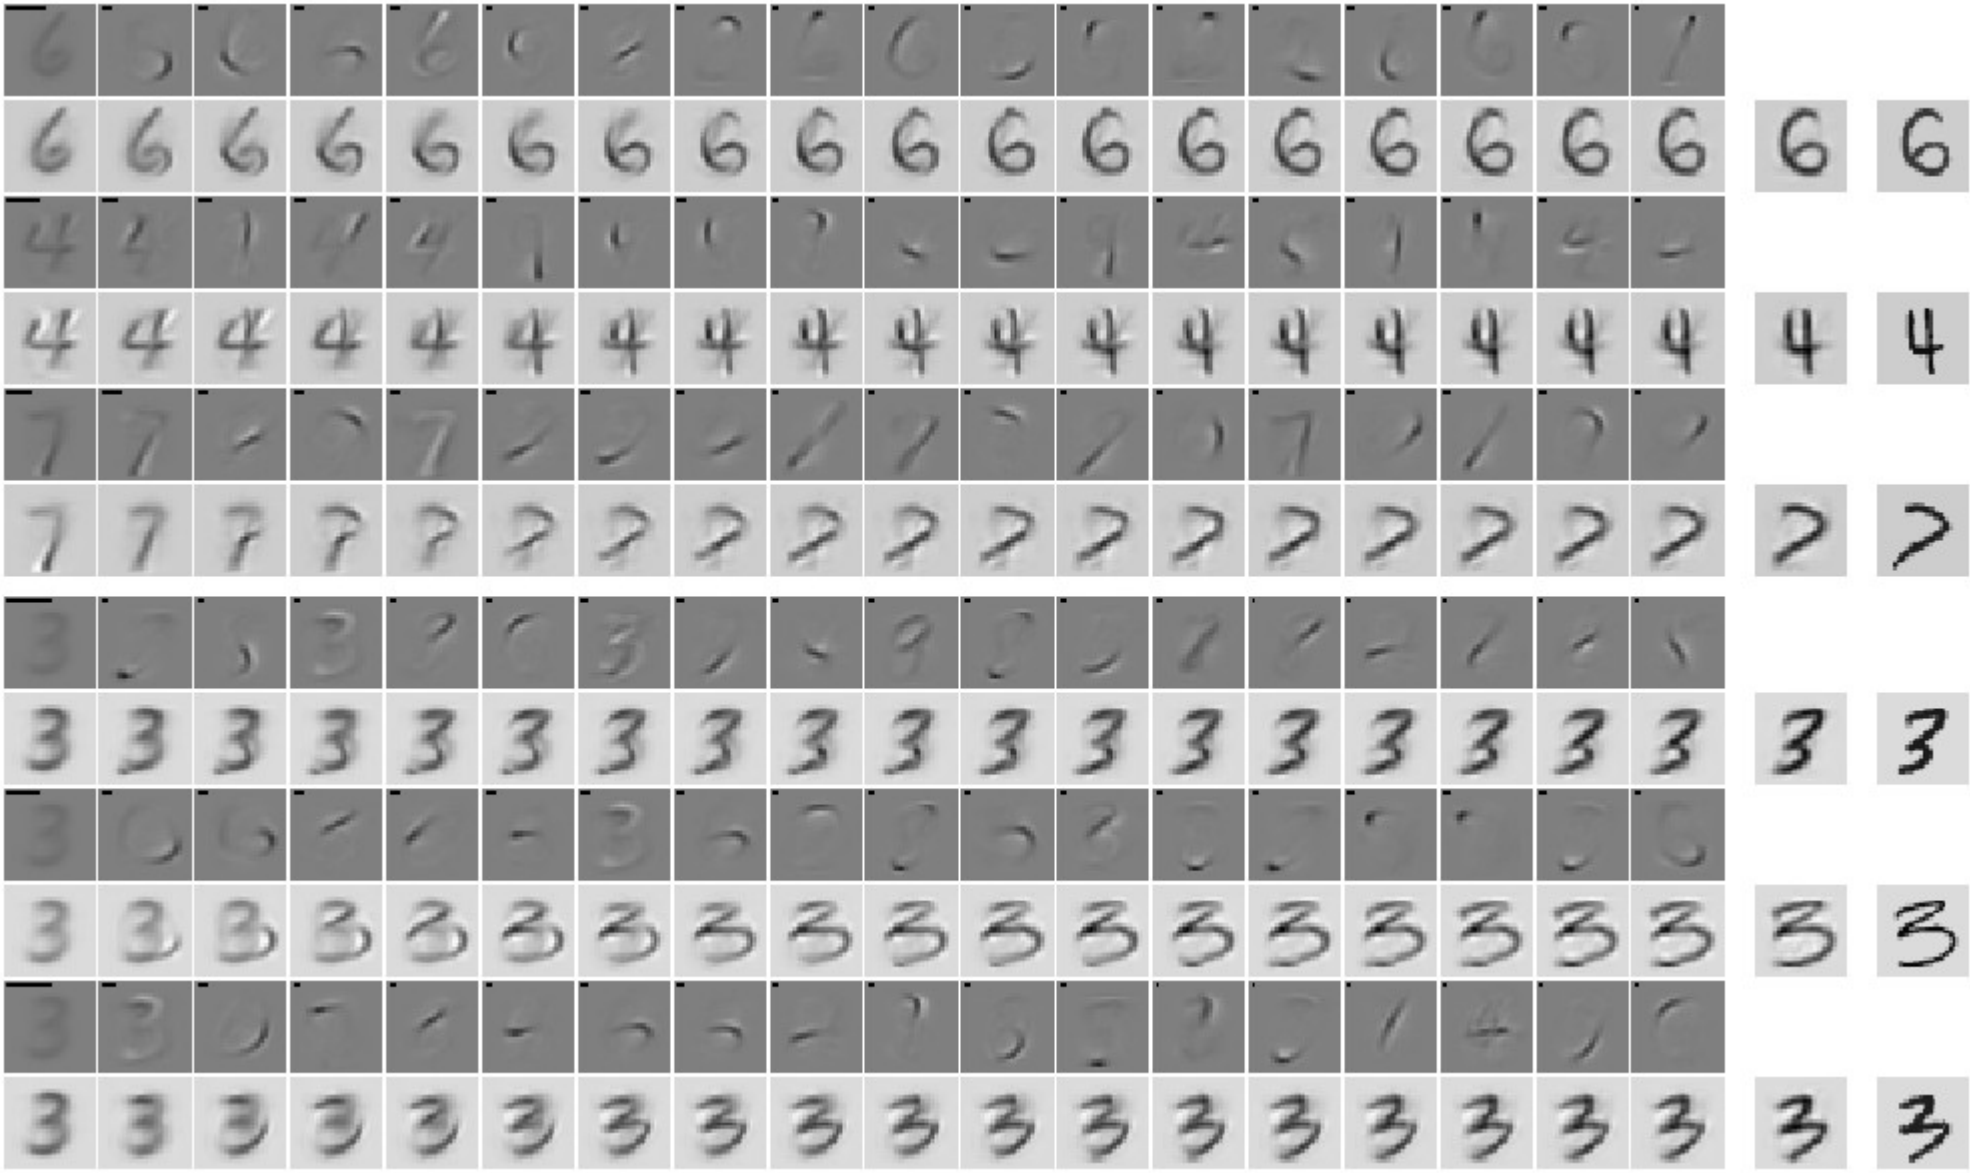

Convolutional Sparse Coding

- **Replace the dot products with dictionary element by convolutions.**
  - Input Y is a full image
  - Each code component Zk is a feature map (an image)
  - Each dictionary element is a convolution kernel
- **Regular sparse coding** $E(Y,Z) = ||Y - \sum_k W_k Z_k||^2 + \alpha \sum_k |Z_k|$
- **Convolutional S.C.** $E(Y,Z) = ||Y - \sum_k W_k * Z_k||^2 + \alpha \sum_k |Z_k|$

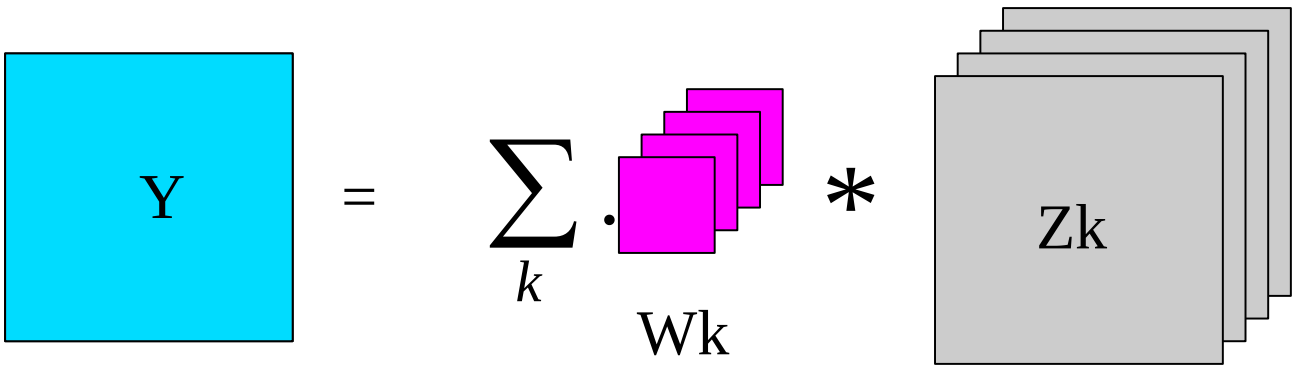

"deconvolutional networks" [Zeiler, Taylor, Fergus CVPR 2010]

Convolutional PSD: Encoder with a soft sh() Function

Convolutional Formulation

- Extend sparse coding from ***PATCH*** to ***IMAGE***

$$\mathcal{L}(x, z, \mathcal{D}) = \frac{1}{2}||x - \sum_{k=1}^{K} \mathcal{D}_k * z_k||_2^2 + \sum_{k=1}^{K} ||z_k - f(W^k * x)||_2^2 + |z|_1$$

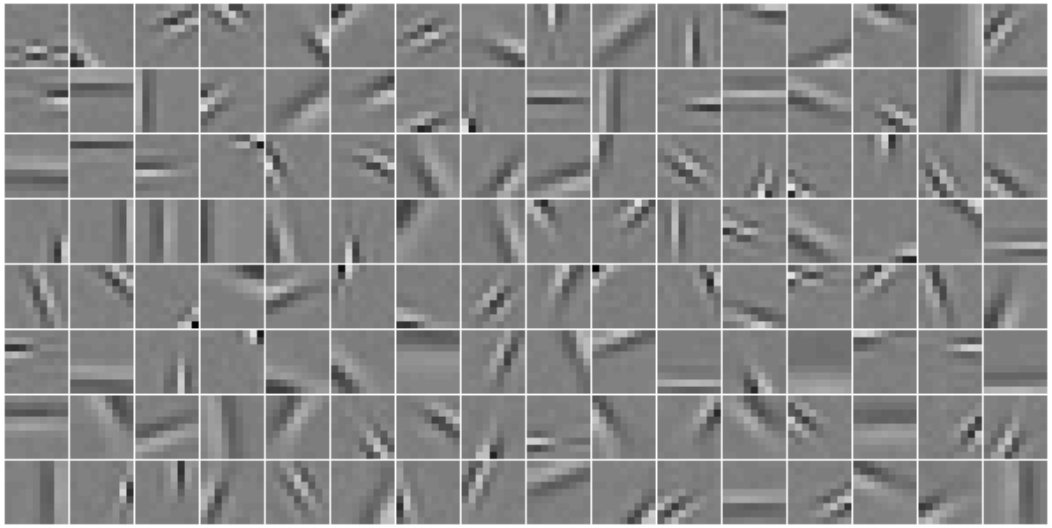

- ***PATCH*** based learning

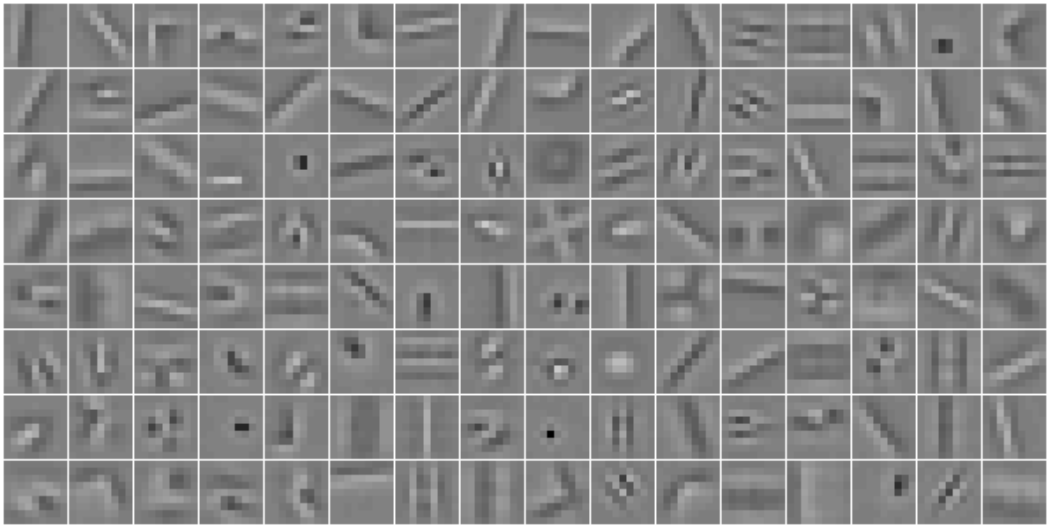

- ***CONVOLUTIONAL*** learning

Convolutional Sparse Auto-Encoder on Natural Images

- Filters and Basis Functions obtained with 1, 2, 4, 8, 16, 32, and 64 filters.

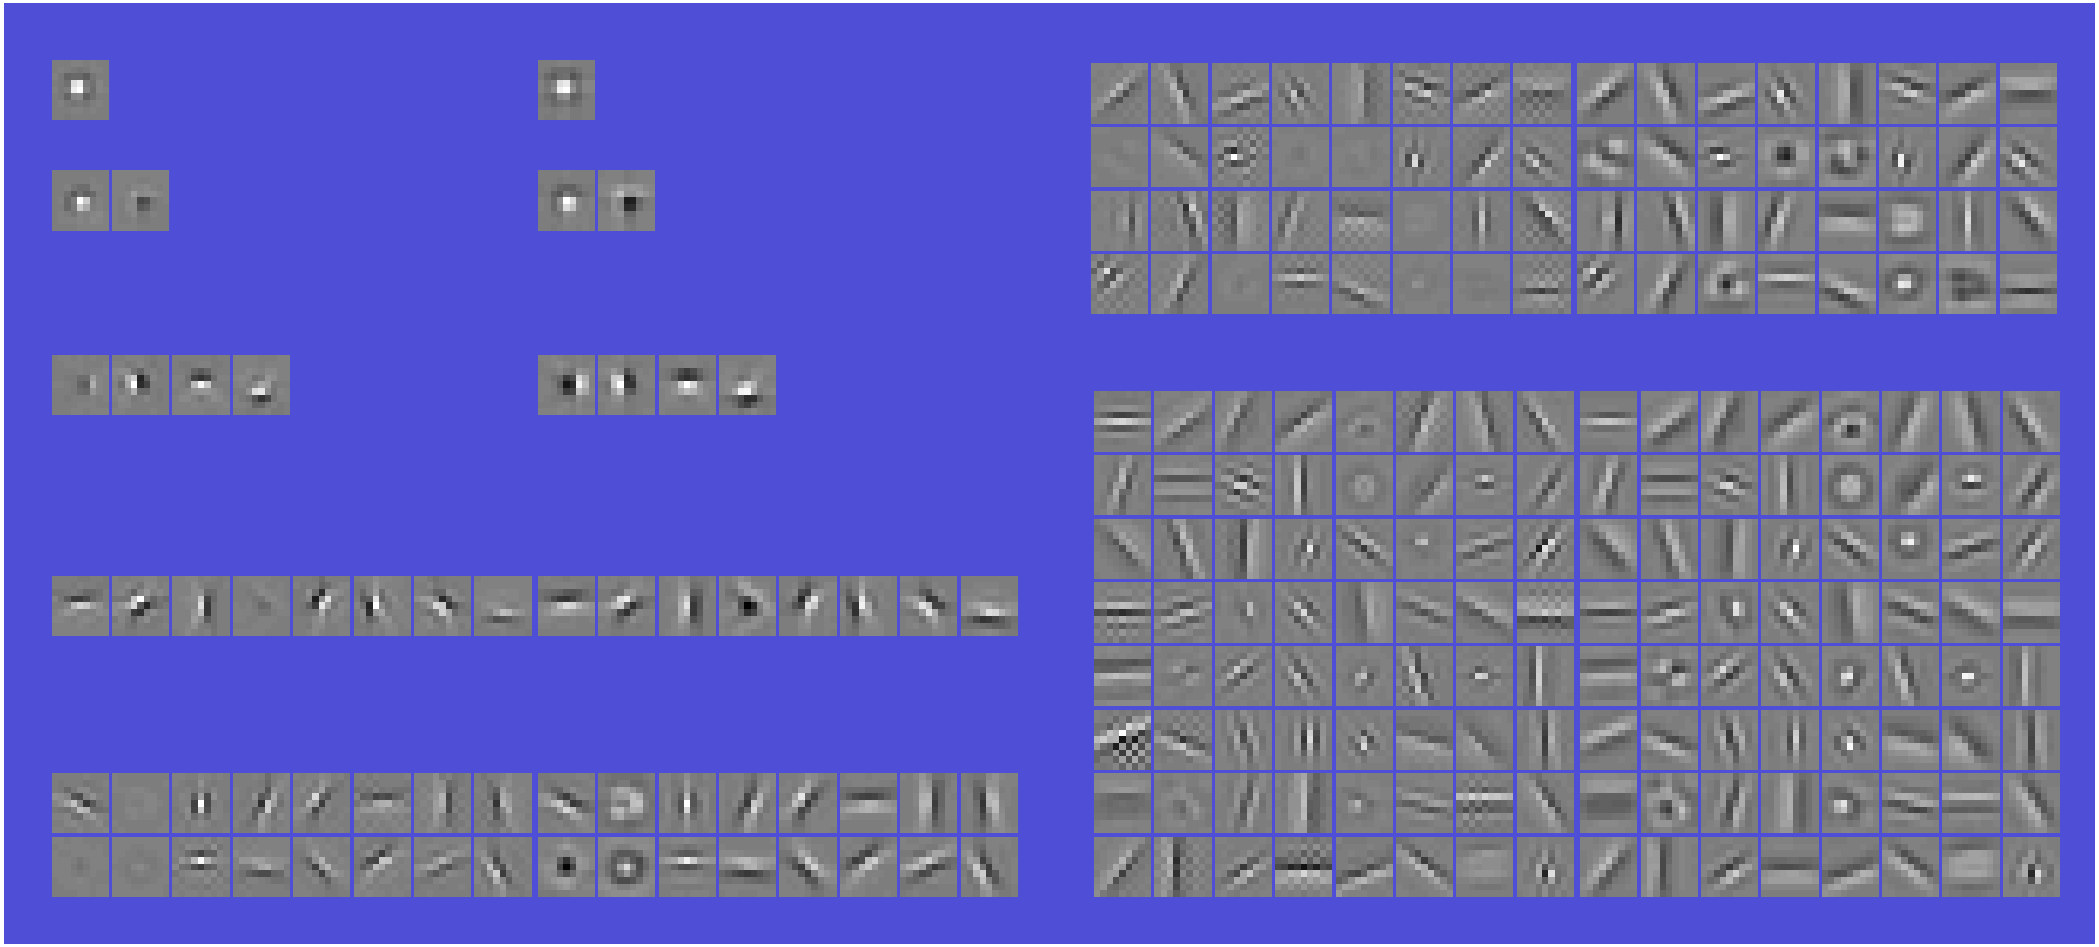

Using PSD to Train a Hierarchy of Features

- Phase 1: train first layer using PSD

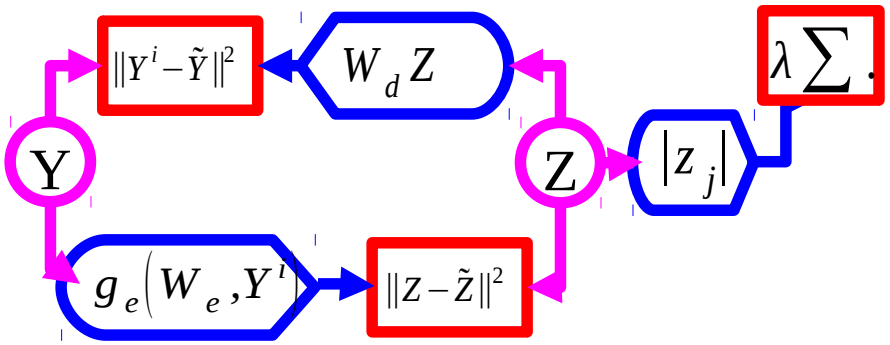

FEATURES

Using PSD to Train a Hierarchy of Features

- Phase 1: train first layer using PSD
- Phase 2: use encoder + absolute value as feature extractor

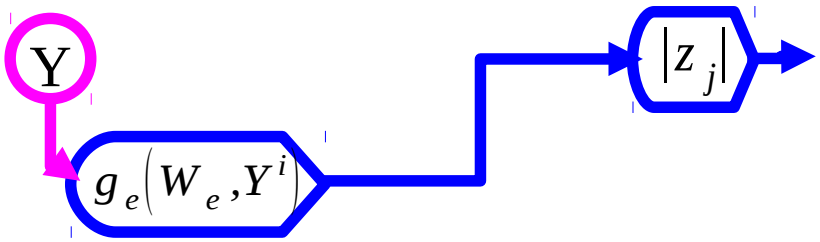

FEATURES

Using PSD to Train a Hierarchy of Features

- Phase 1: train first layer using PSD
- Phase 2: use encoder + absolute value as feature extractor
- Phase 3: train the second layer using PSD

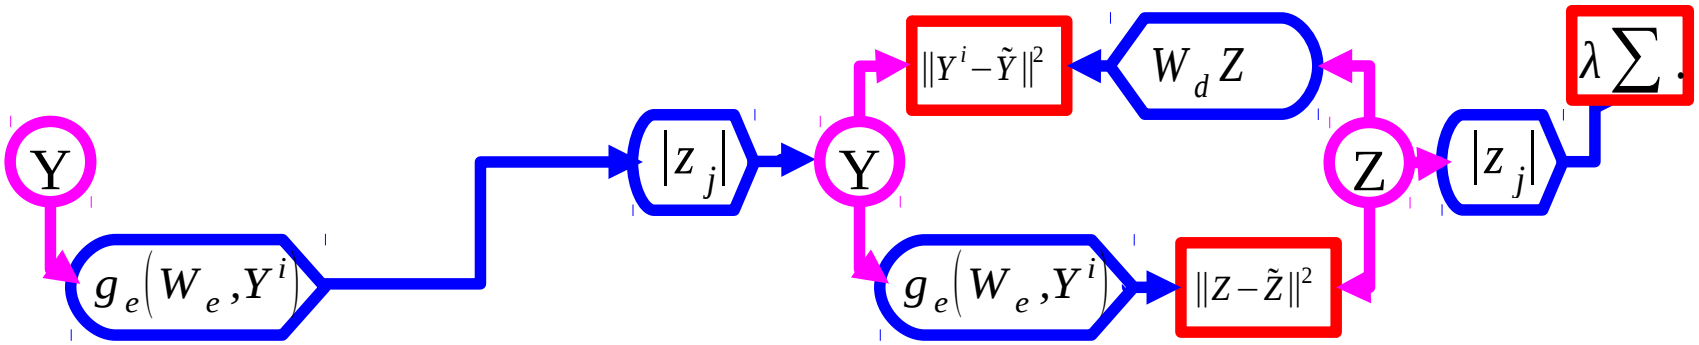

FEATURES

Using PSD to Train a Hierarchy of Features

- Phase 1: train first layer using PSD
- Phase 2: use encoder + absolute value as feature extractor
- Phase 3: train the second layer using PSD
- Phase 4: use encoder + absolute value as 2nd feature extractor

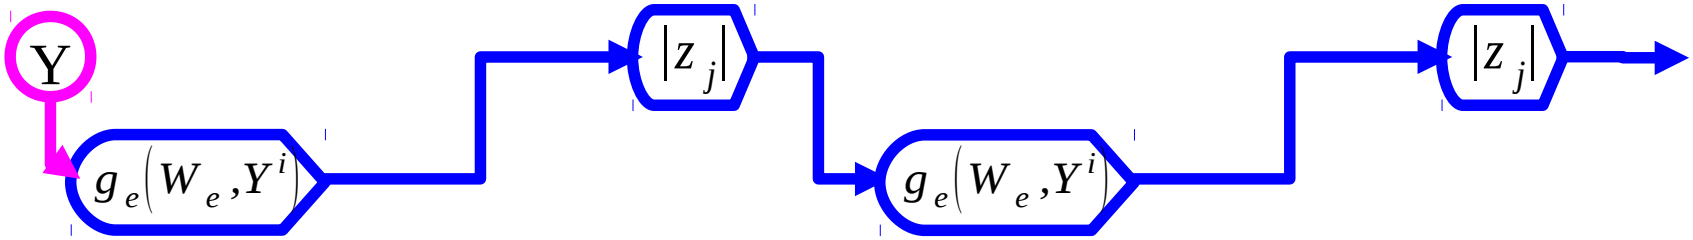

Using PSD to Train a Hierarchy of Features

- Phase 1: train first layer using PSD
- Phase 2: use encoder + absolute value as feature extractor
- Phase 3: train the second layer using PSD
- Phase 4: use encoder + absolute value as 2nd feature extractor
- Phase 5: train a supervised classifier on top
- Phase 6 (optional): train the entire system with supervised back-propagation

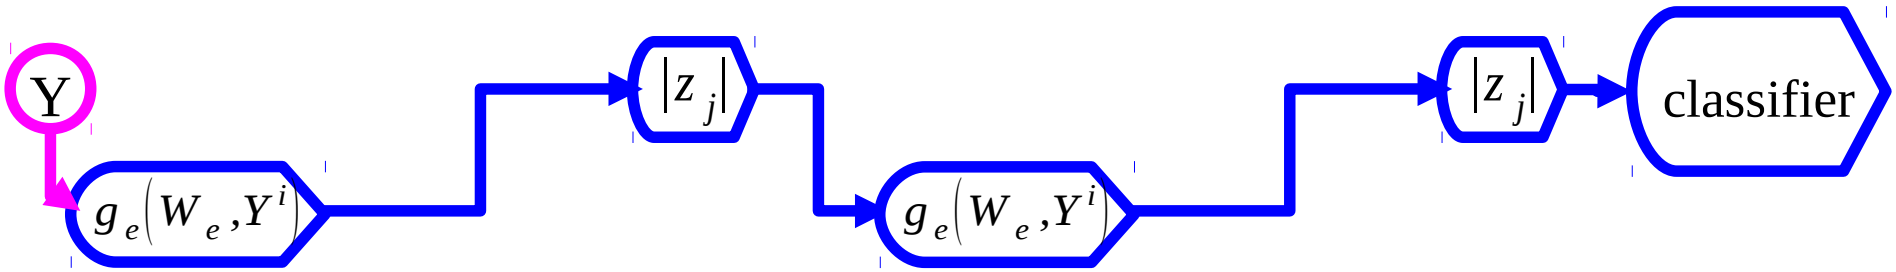

Unsupervised Learning:
Invariant Features

Learning Invariant Features with L2 Group Sparsity

- Unsupervised PSD ignores the spatial pooling step.
- Could we devise a similar method that learns the pooling layer as well?
- Idea [Hyvarinen & Hoyer 2001]: group sparsity on pools of features
  - Minimum number of pools must be non-zero
  - Number of features that are on within a pool doesn't matter
  - Pools tend to regroup similar features

$$E(Y,Z) = \|Y - W_d Z\|^2 + \|Z - g_e(W_e, Y)\|^2 + \sum_j \sqrt{\sum_{k \in P_j} Z_k^2}$$

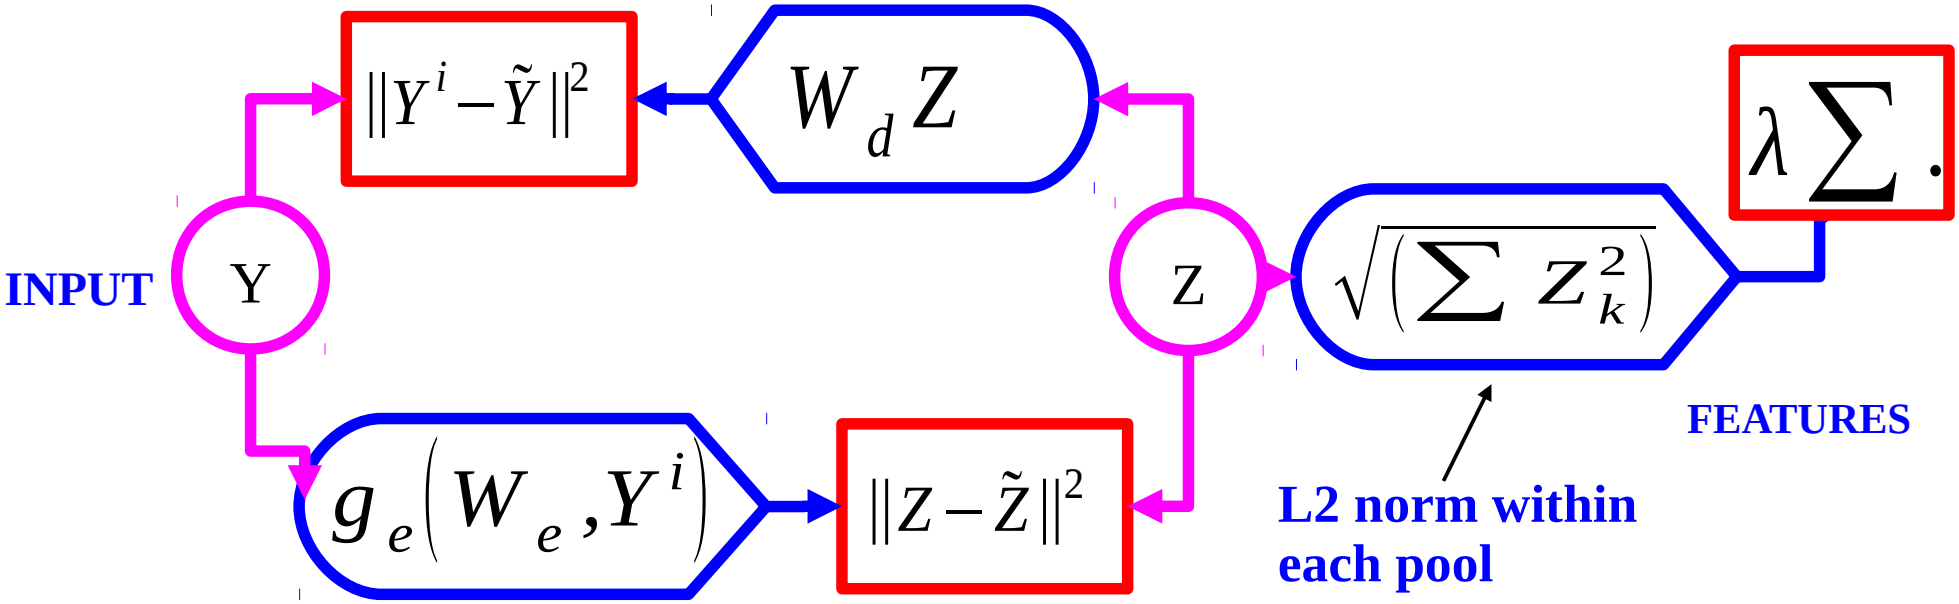

Learning Invariant Features with L2 Group Sparsity

- **Idea: features are pooled in group.**
  - Sparsity: sum over groups of L2 norm of activity in group.
- **[Hyvärinen Hoyer 2001]: "subspace ICA"**
  - decoder only, square
- **[Welling, Hinton, Osindero NIPS 2002]: pooled product of experts**
  - encoder only, overcomplete, log student-T penalty on L2 pooling
- **[Kavukcuoglu, Ranzato, Fergus LeCun, CVPR 2010]: Invariant PSD**
  - encoder-decoder (like PSD), overcomplete, L2 pooling
- **[Le et al. NIPS 2011]: Reconstruction ICA**
  - Same as [Kavukcuoglu 2010] with linear encoder and tied decoder
- **[Gregor & LeCun arXiv:1006:0448, 2010] [Le et al. ICML 2012]**
  - Locally-connect non shared (tiled) encoder-decoder

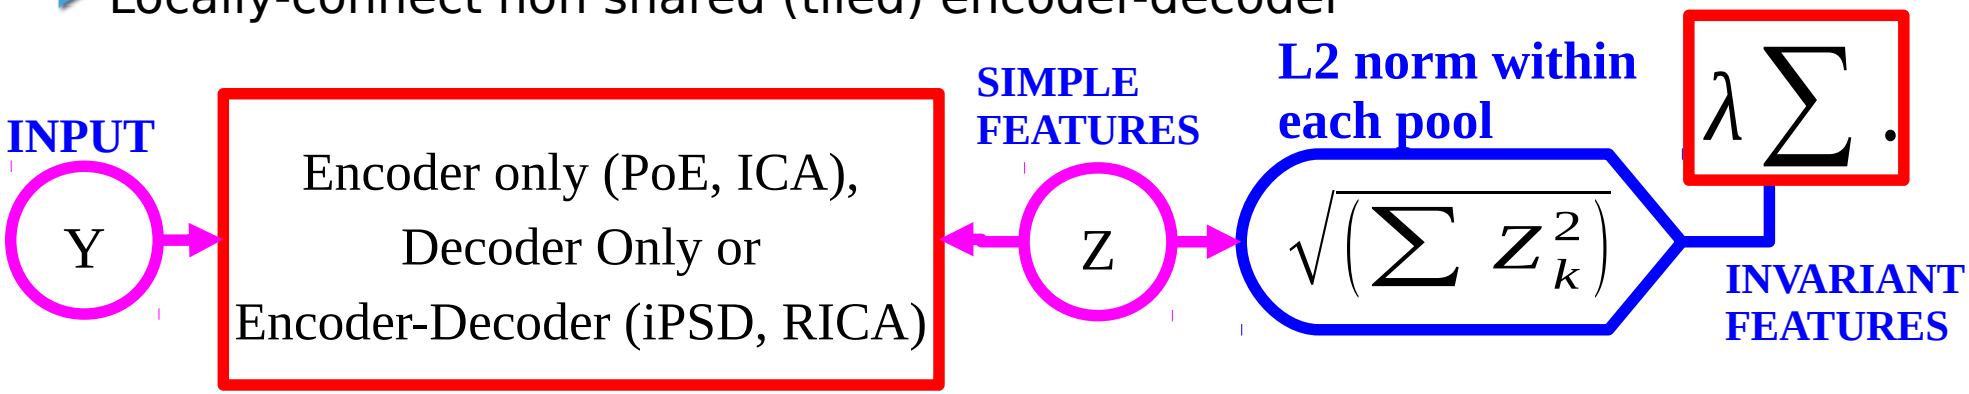

Groups are local in a 2D Topographic Map

- The filters arrange themselves spontaneously so that similar filters enter the same pool.
- The pooling units can be seen as complex cells
- Outputs of pooling units are invariant to local transformations of the input
  - For some it's translations, for others rotations, or other transformations.

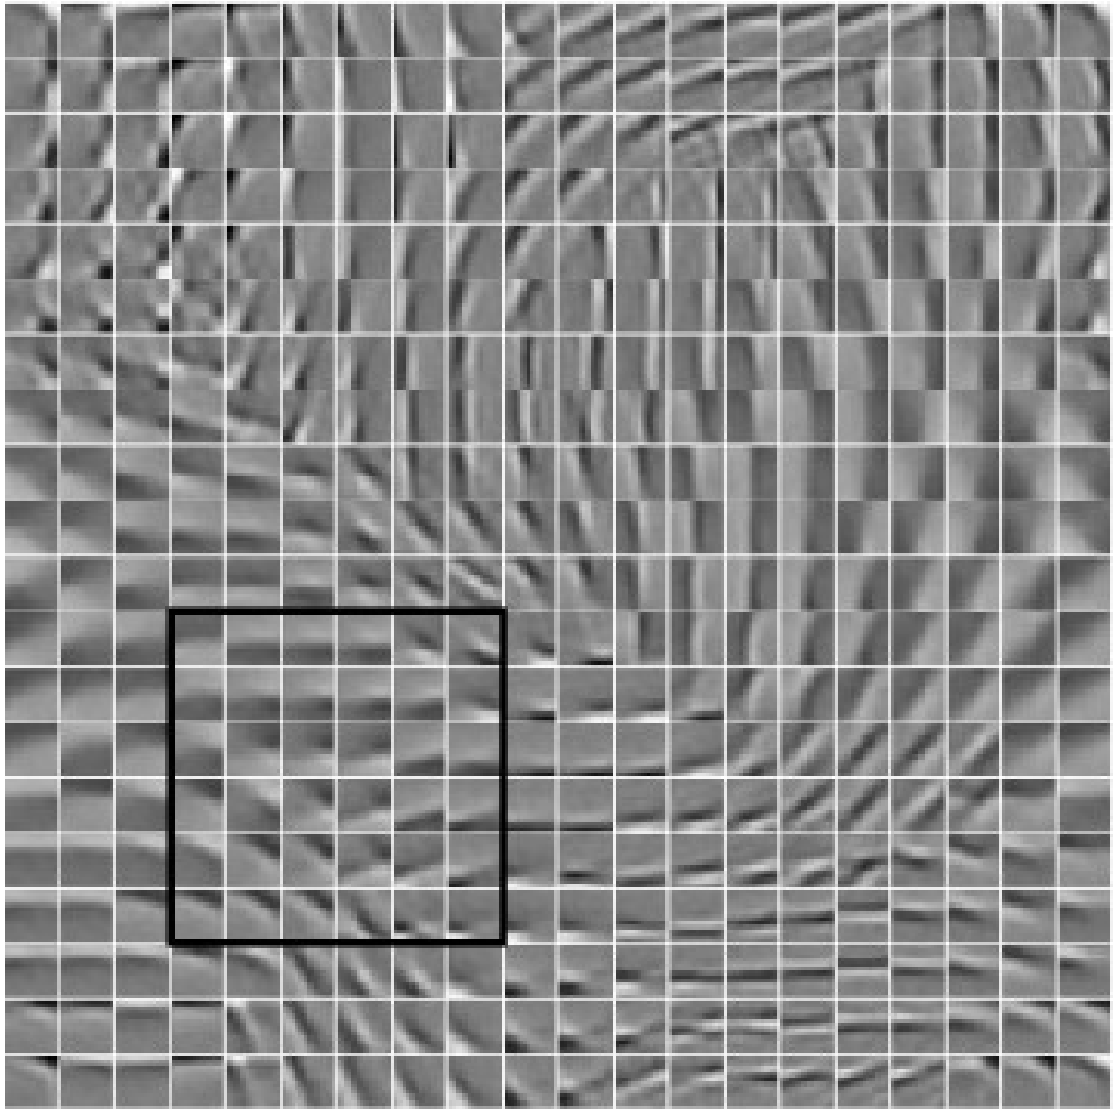

Image-level training, local filters but no weight sharing

- Training on 115x115 images. Kernels are 15x15 (not shared across space!)
  - [Gregor & LeCun 2010]
  - Local receptive fields
  - No shared weights
  - 4x overcomplete
  - L2 pooling
  - Group sparsity over pools

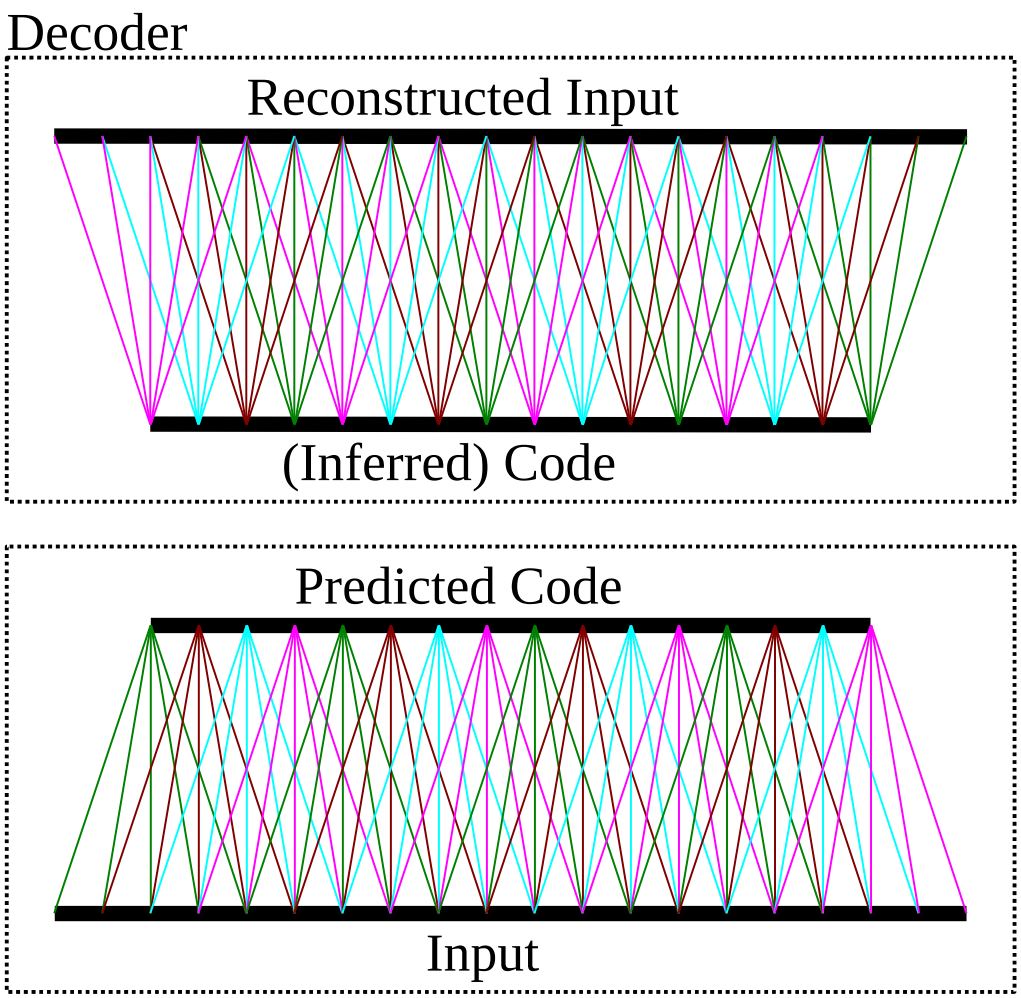

Image-level training, local filters but no weight sharing

- Training on 115x115 images. Kernels are 15x15 (not shared across space!)

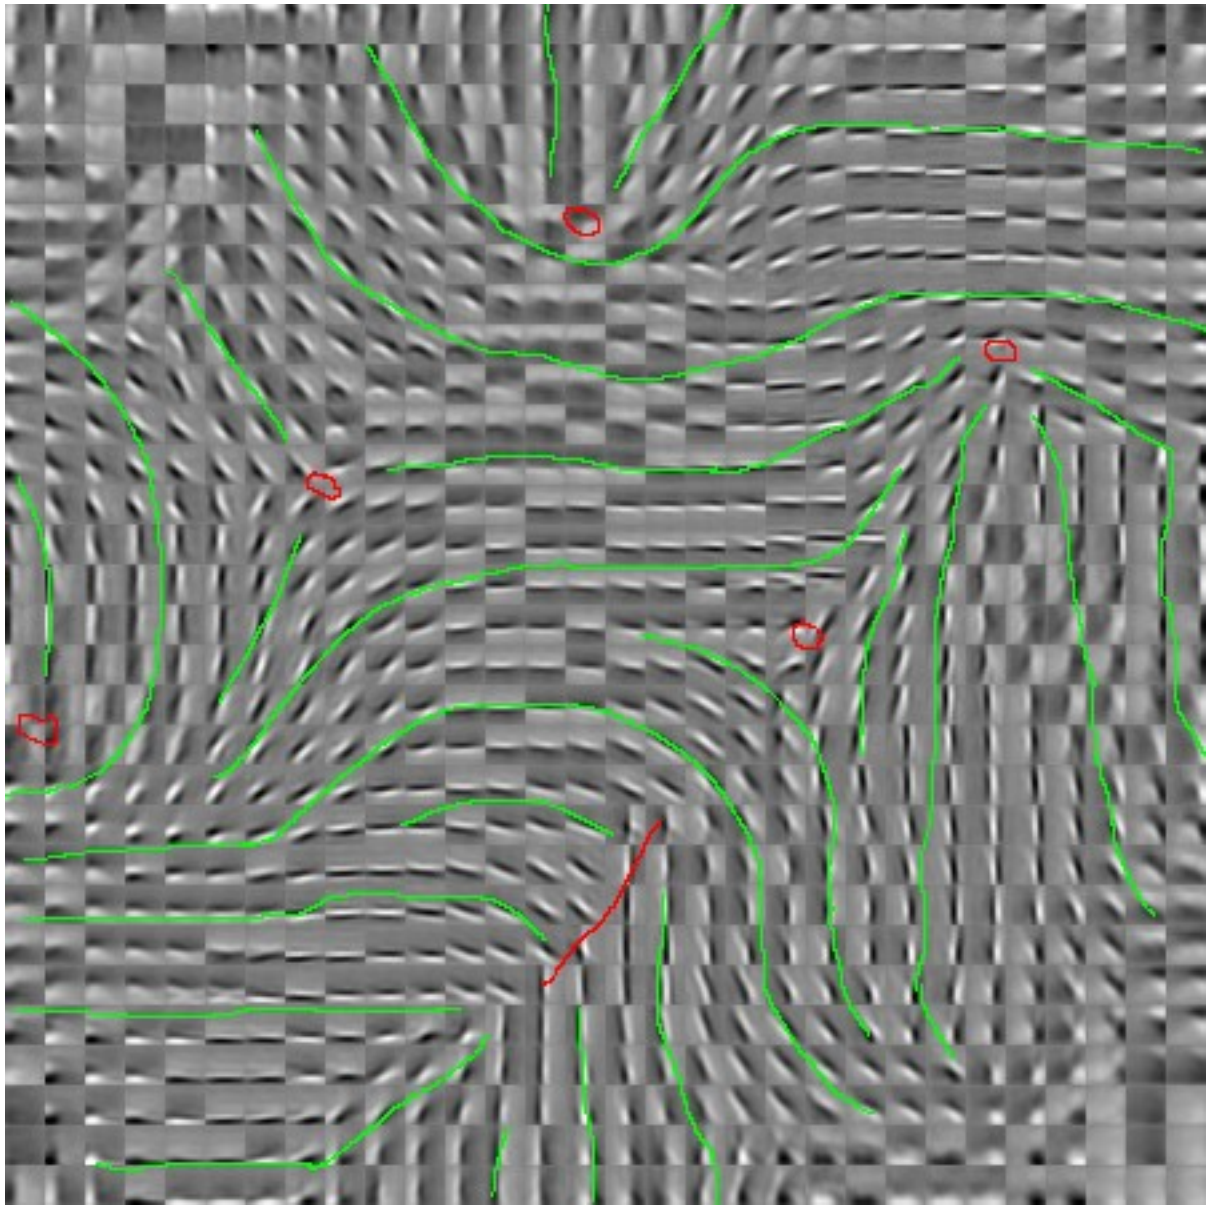

Topographic Maps

K Obermayer and GG Blasdel, Journal of Neuroscience, Vol 13, 4114-4129 (**Monkey**)

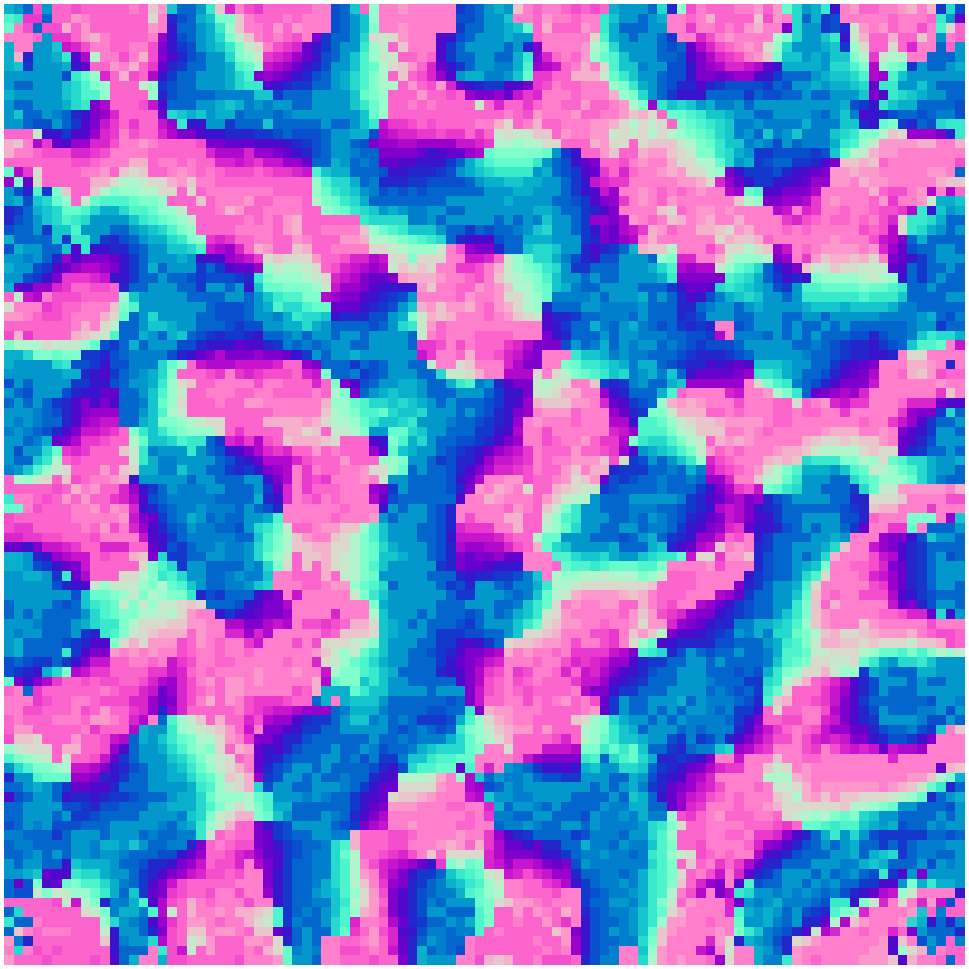

119x119 Image Input

100x100 Code

20x20 Receptive field size

sigma=5

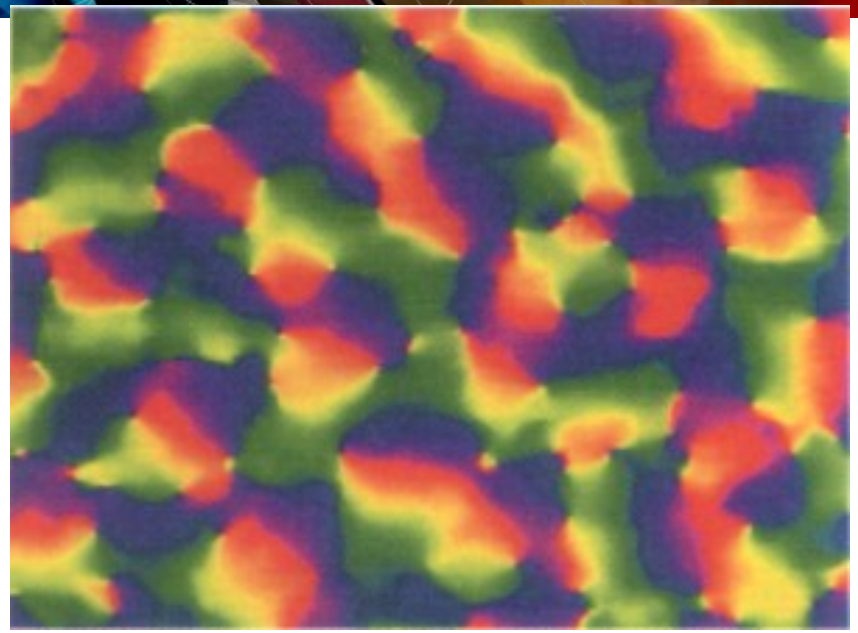

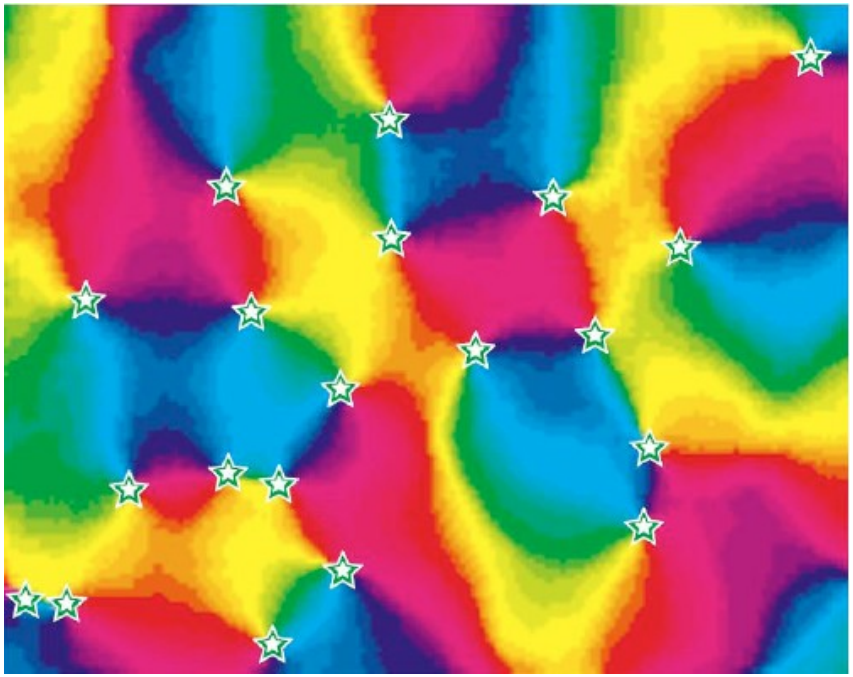

Michael C. Crair, et. al. The Journal of Neurophysiology Vol. 77 No. 6 June 1997, pp. 3381-3385 (**Cat**)

Image-level training, local filters but no weight sharing

- Color indicates orientation (by fitting Gabors)

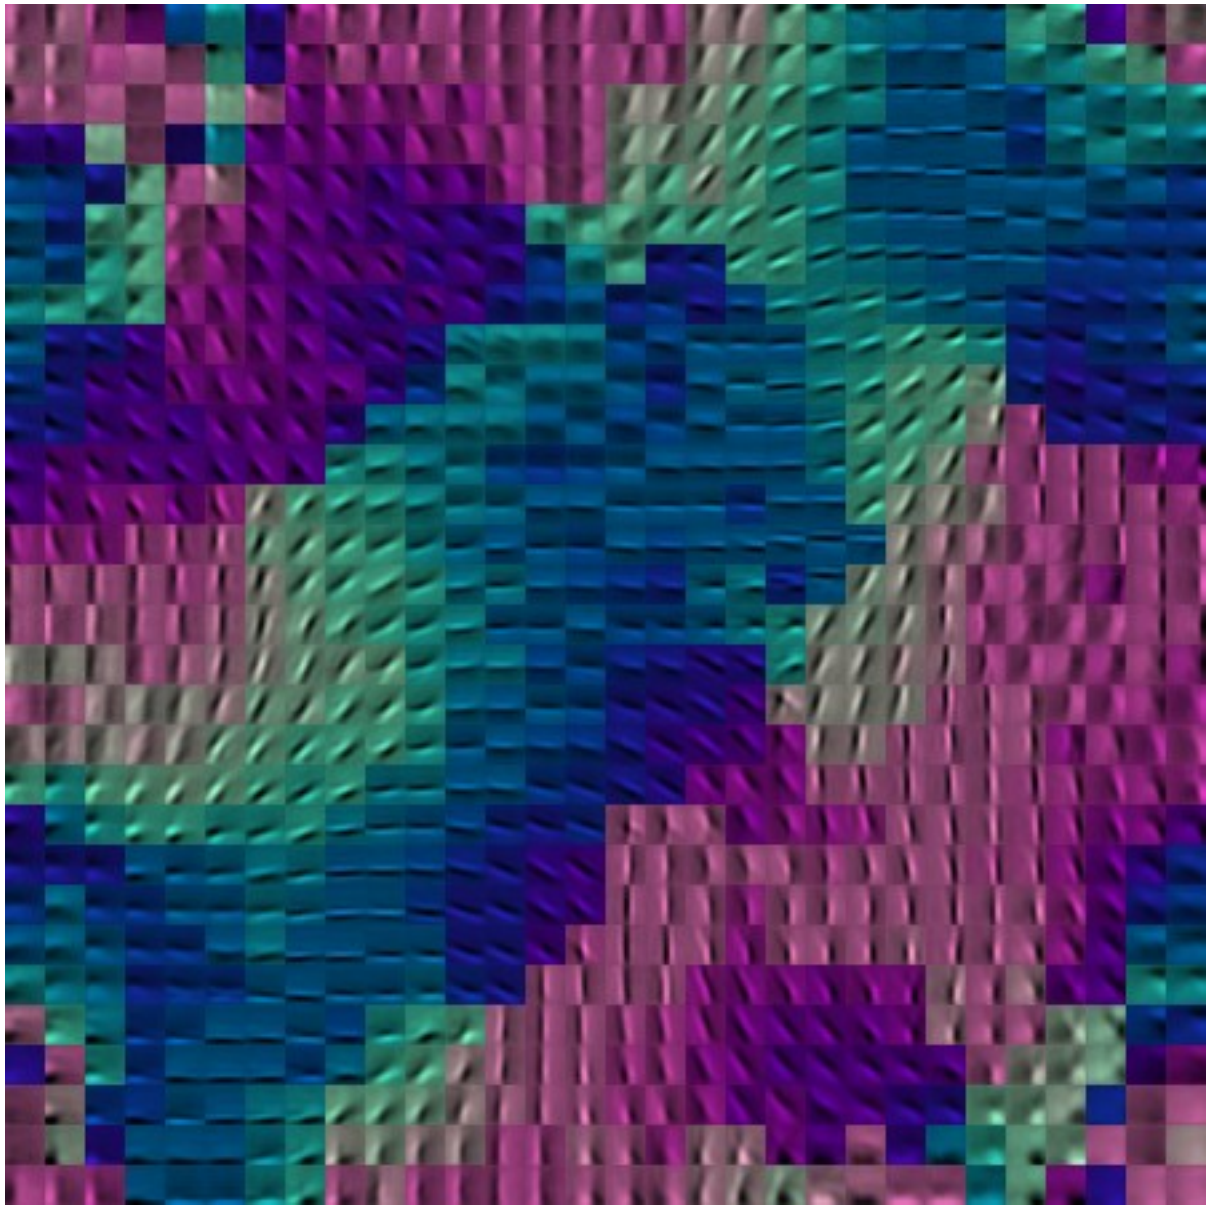

Invariant Features Lateral Inhibition

- Replace the L1 sparsity term by a lateral inhibition matrix
- Easy way to impose some structure on the sparsity

$$\min_{W,Z} \sum_{x \in X} ||Wz - x||^2 + |z|^T S|z|$$

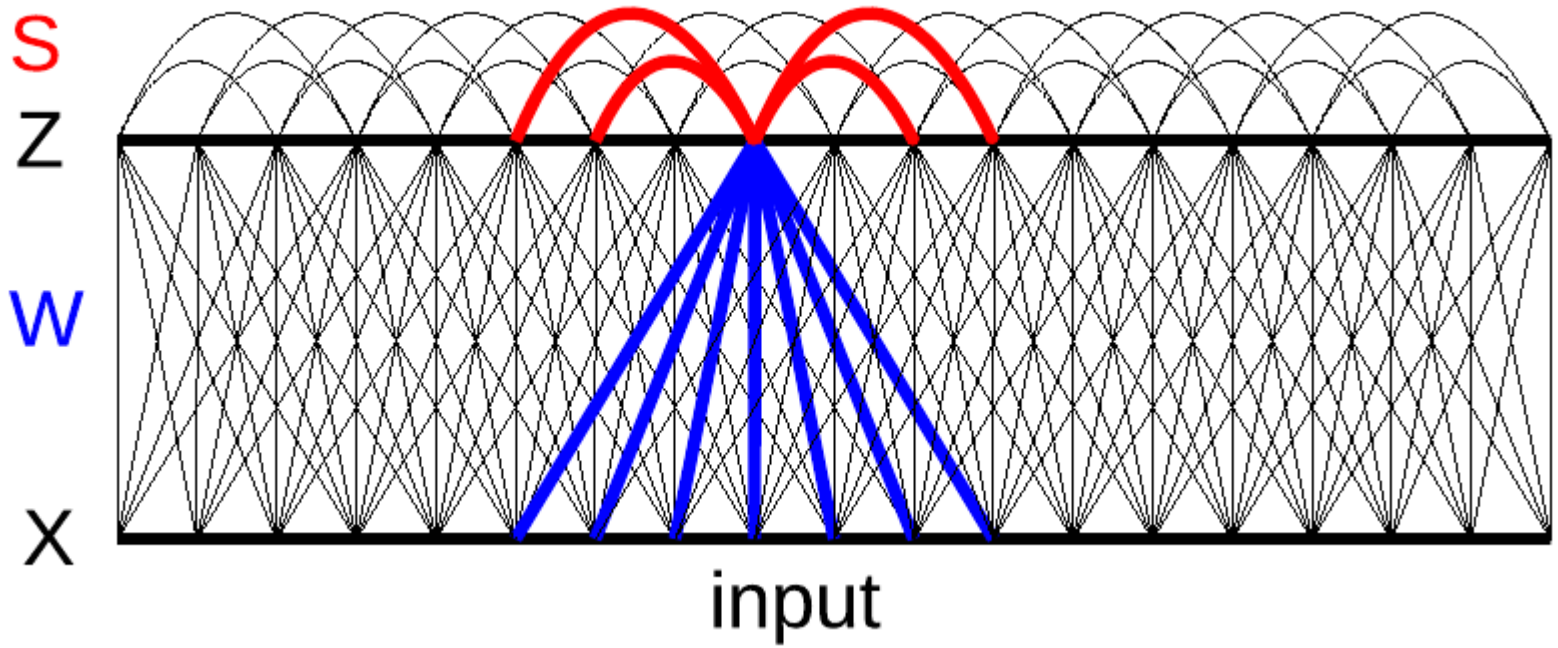

[Gregor, Szlam, LeCun NIPS 2011]

Invariant Features via Lateral Inhibition: Structured Sparsity

- Each edge in the tree indicates a zero in the S matrix (no mutual inhibition)
- Sij is larger if two neurons are far away in the tree

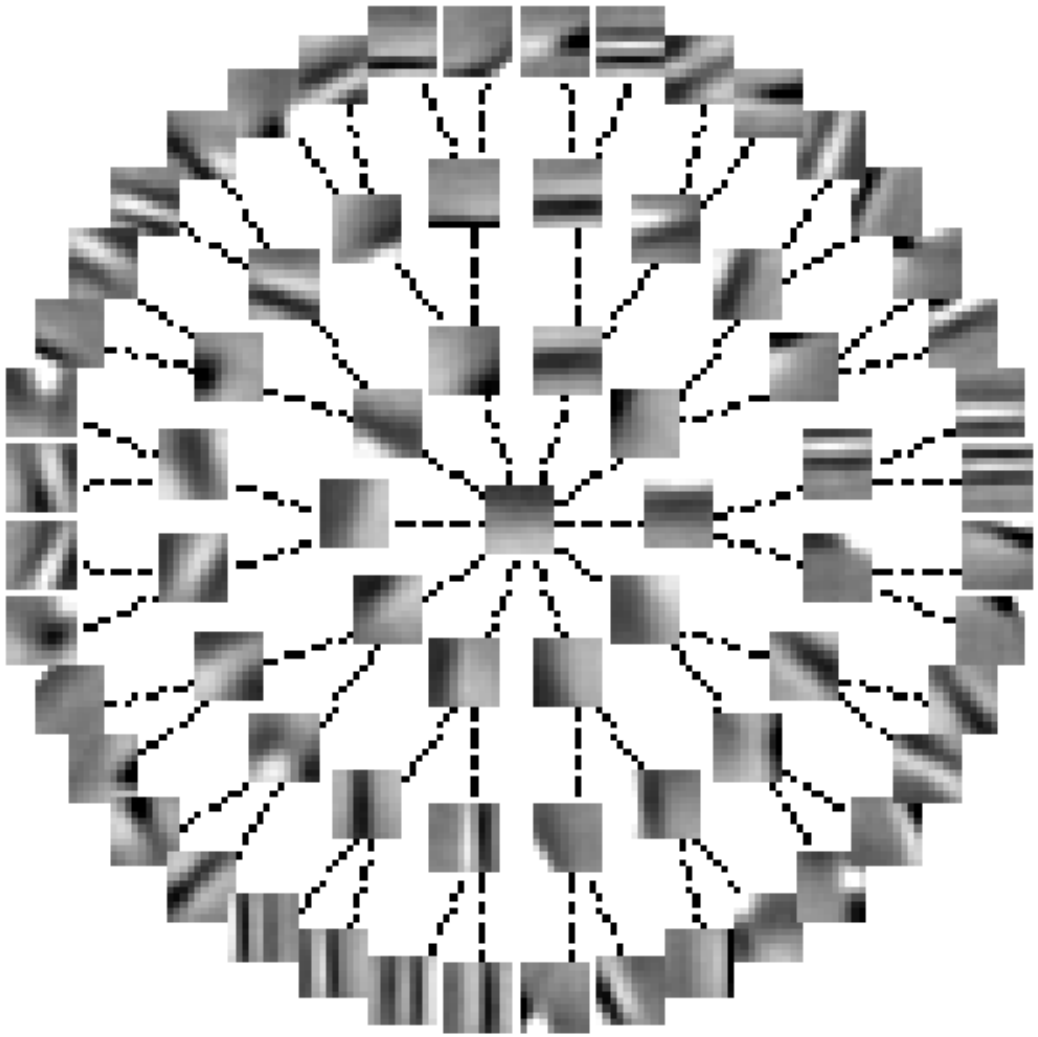

Invariant Features via Lateral Inhibition: Topographic Maps

- **Non-zero values in S form a ring in a 2D topology**
  - Input patches are high-pass filtered

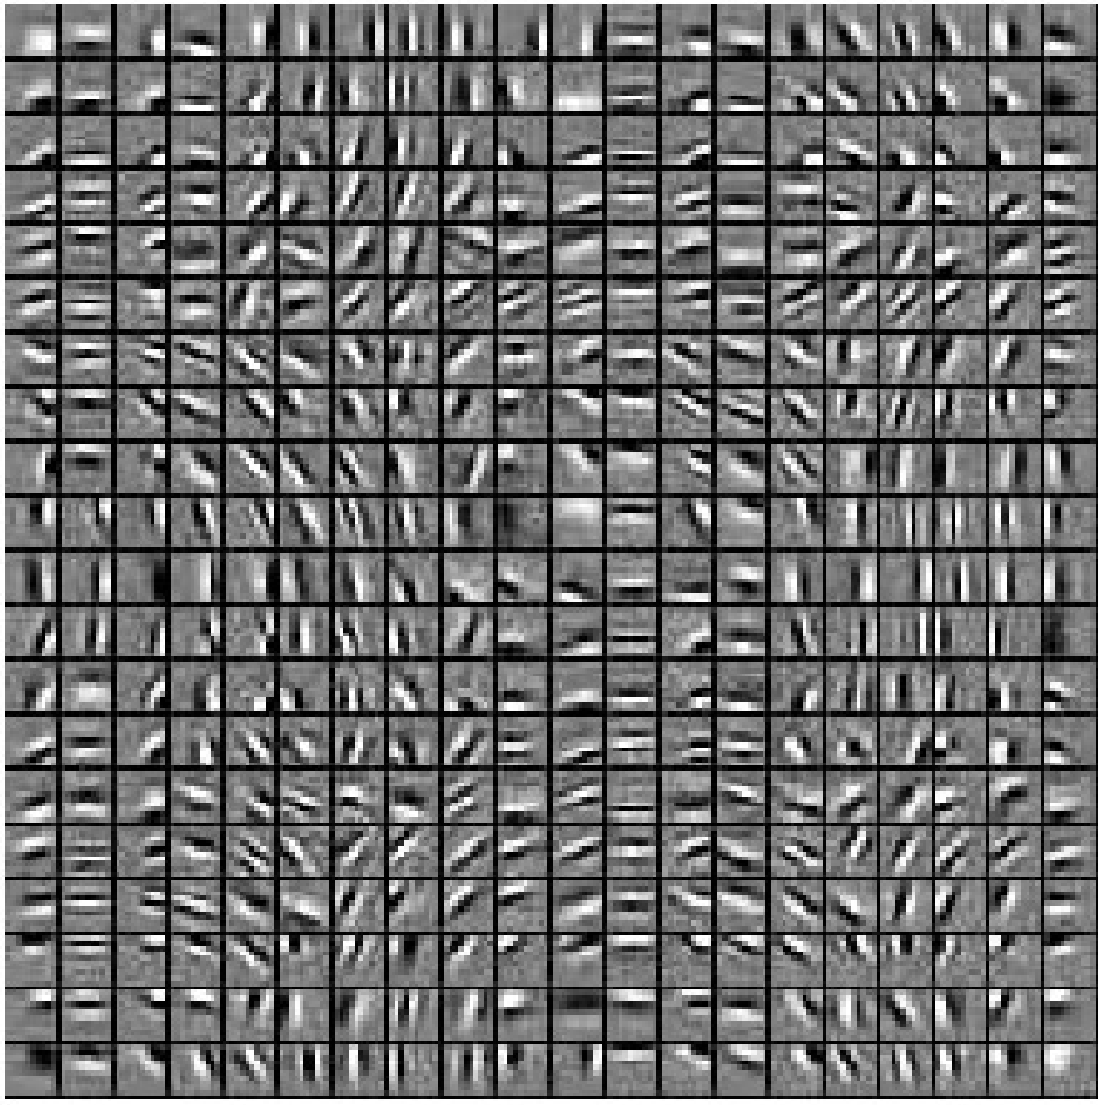

Sparse Auto-Encoder with "Slow Feature" Penalty

Y LeCun

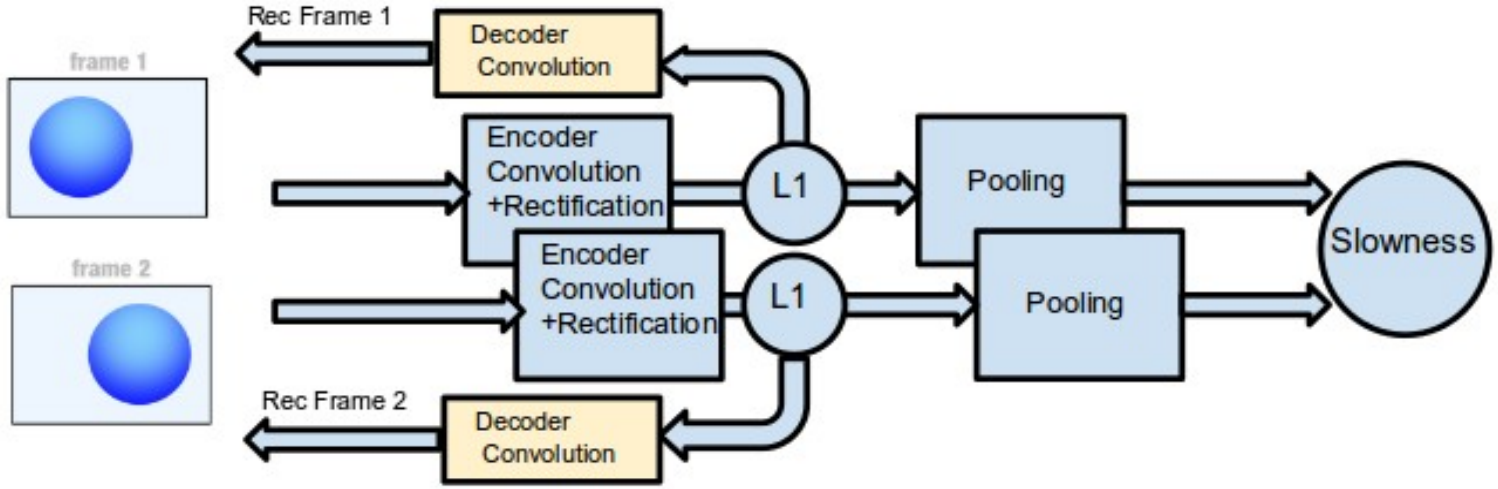

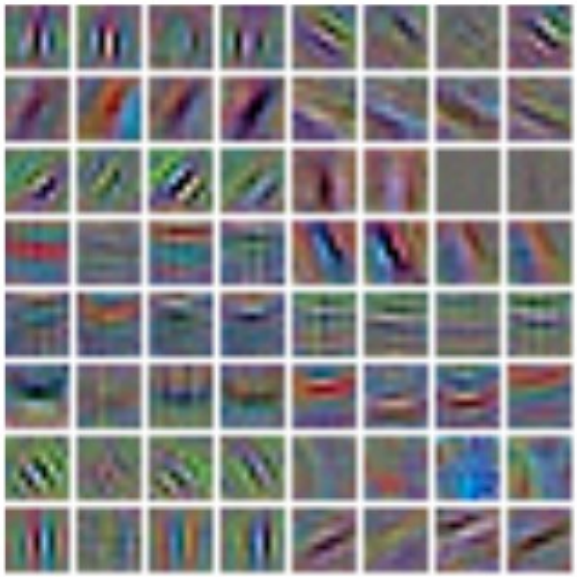

Supervised filters CIFAR10

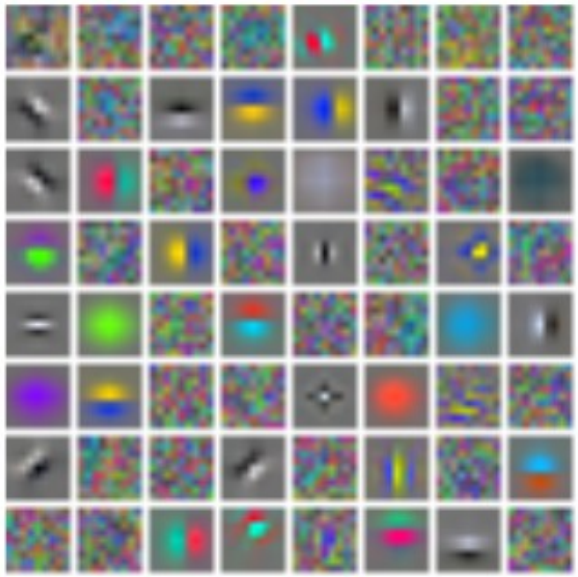

sparse conv. auto-encoder

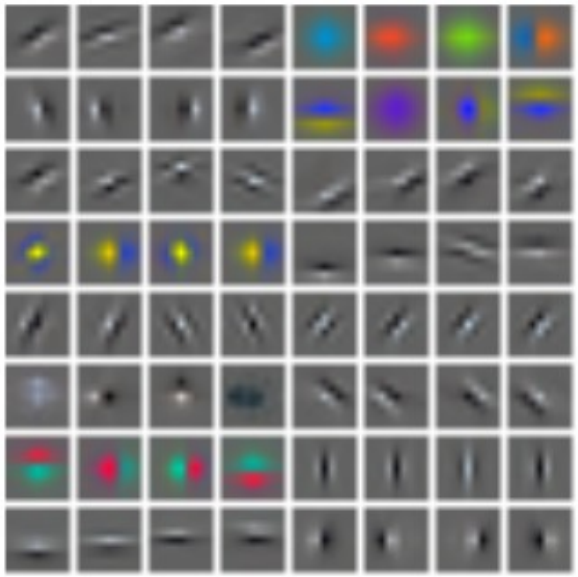

slow & sparse conv. auto-encoder

Trained on YouTube videos

[Goroshin et al. Arxiv:1412.6056]

Invariant Features through Temporal Constancy

- Object is cross-product of object type and instantiation parameters
  - Mapping units [Hinton 1981], capsules [Hinton 2011]

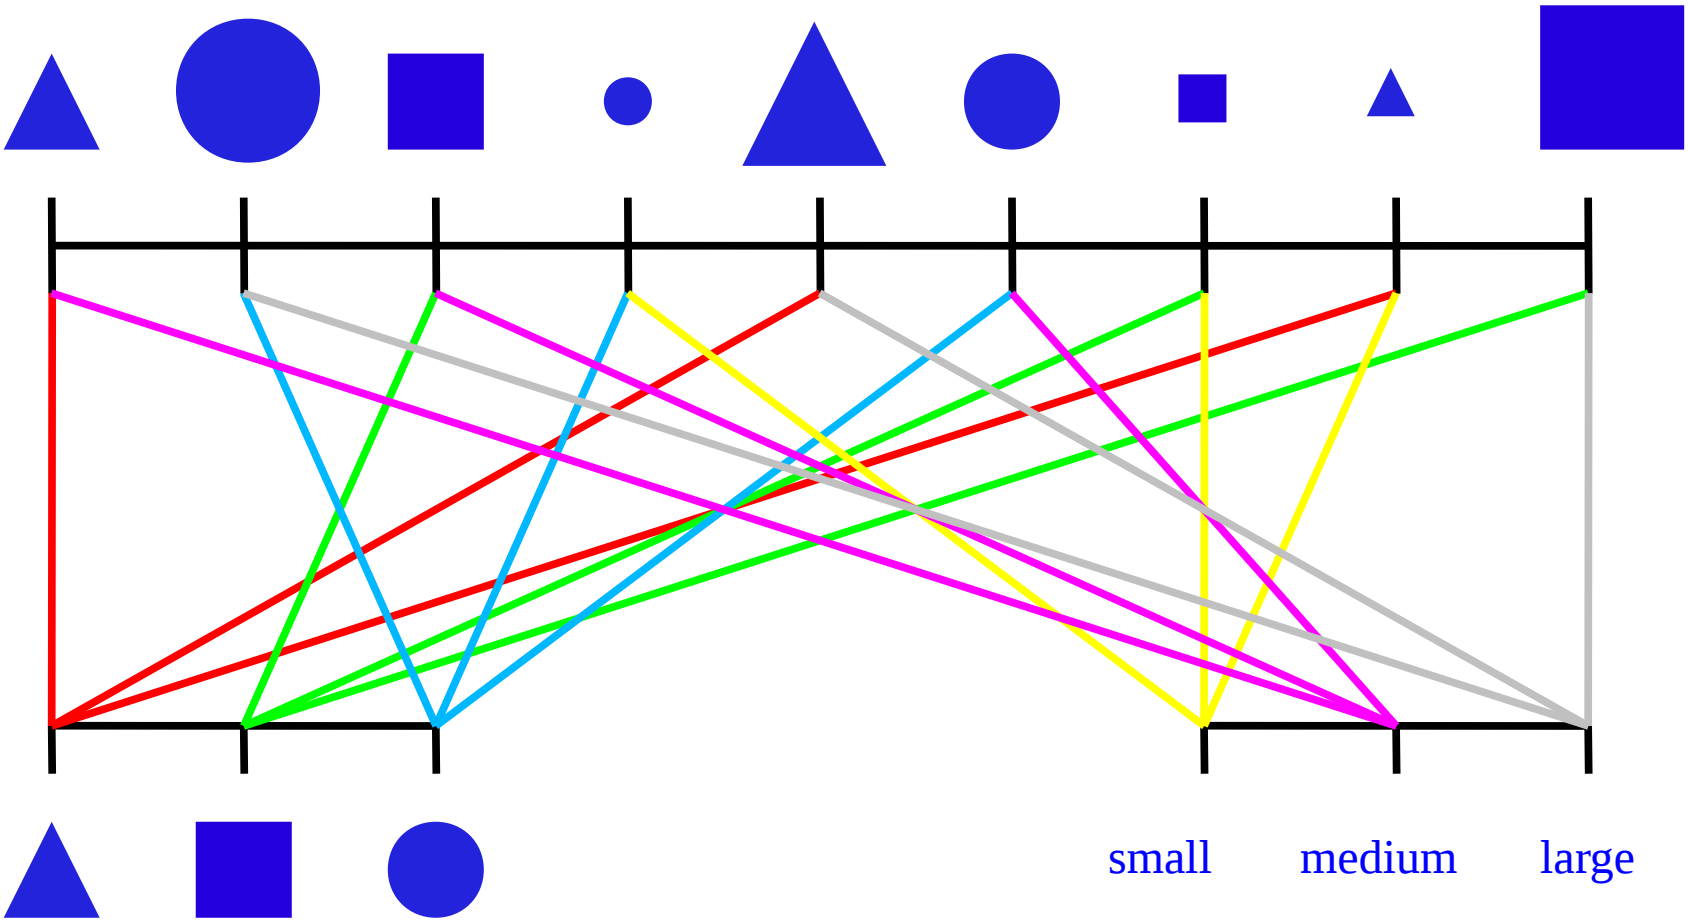

[Karol Gregor et al.]

What-Where Auto-Encoder Architecture

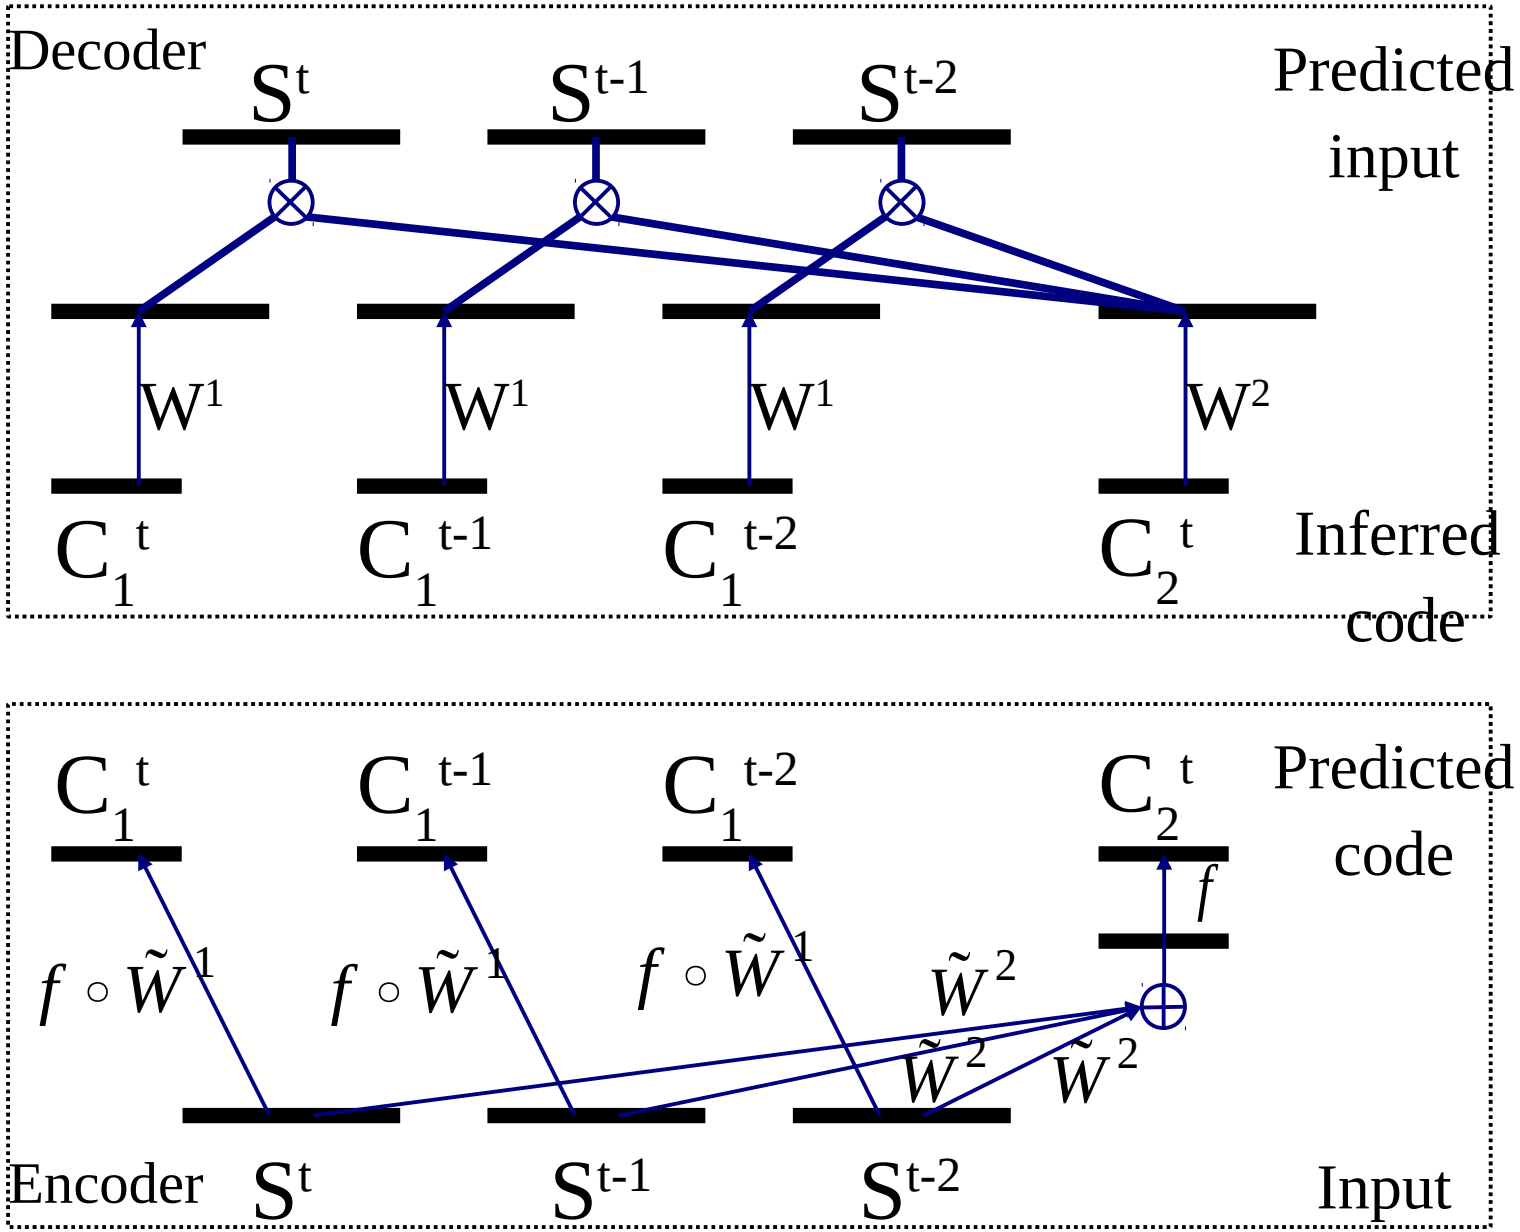

Low-Level Filters Connected to Each Complex Cell

C1
(where)

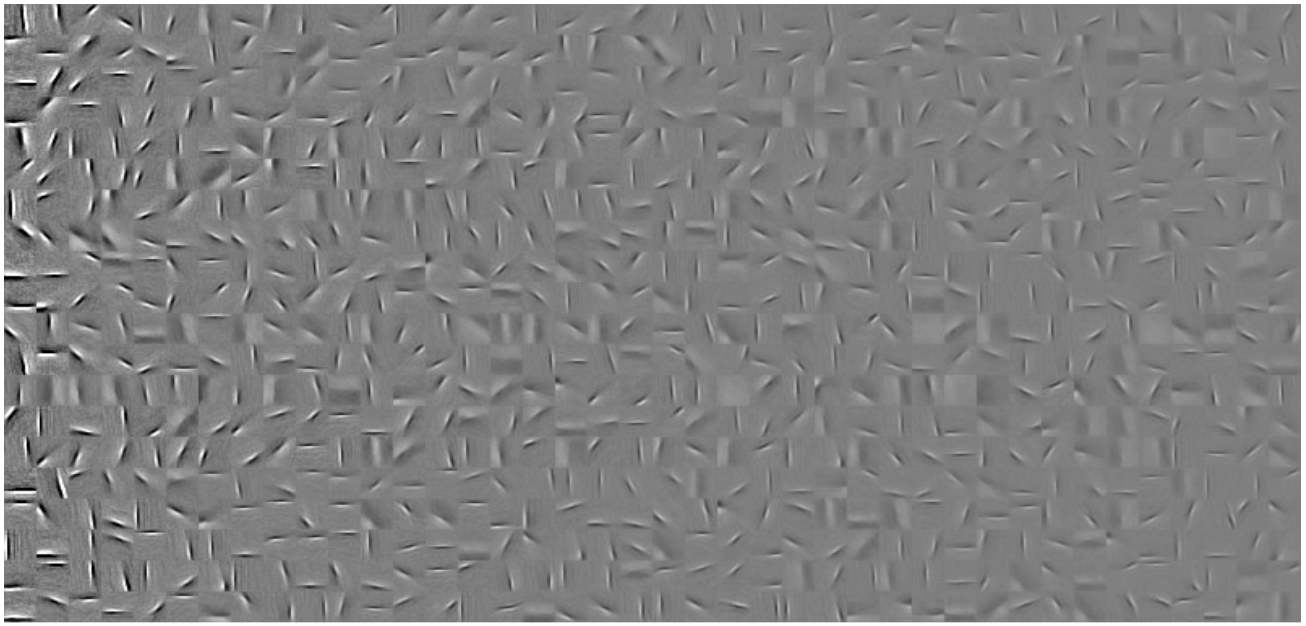

C2
(what)

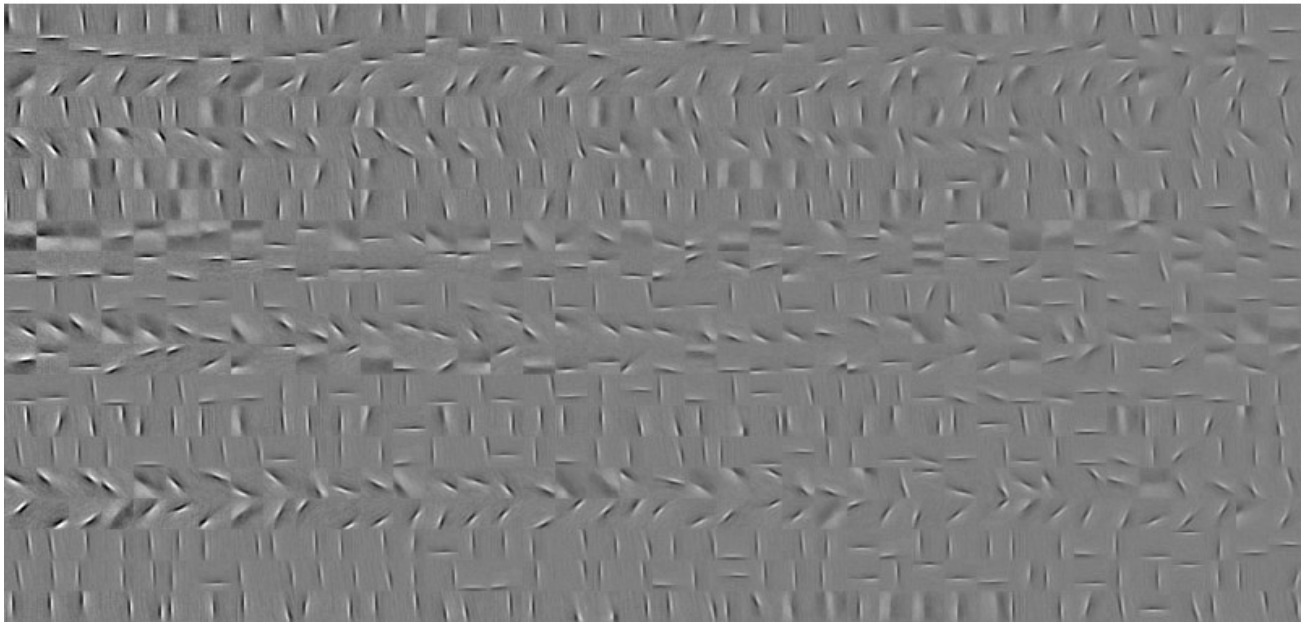

Generating images

Input

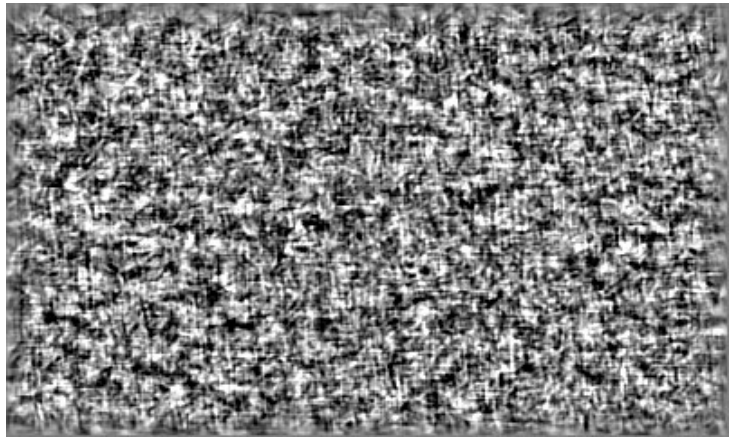

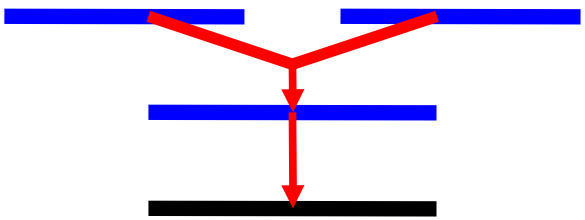

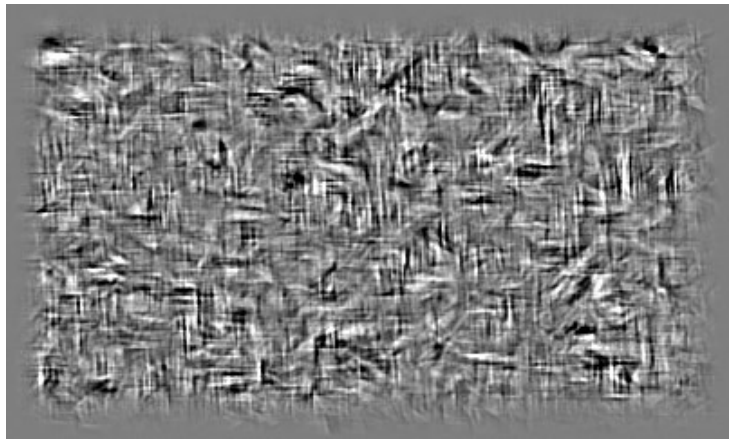

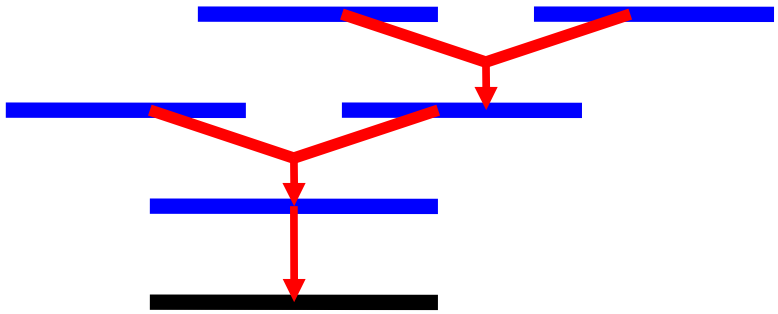

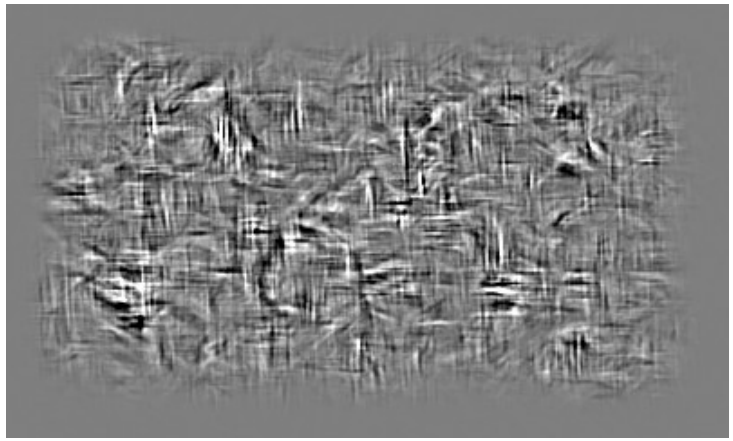

Future
Challenges

Future Challenges

- **Integrated feed-forward and feedback**
  - Deep Boltzmann machine do this, but there are issues of scalability.
- **Integrating supervised and unsupervised learning in a single algorithm**
- **Integrating deep learning and structured prediction ("reasoning")**
  - This has been around since the 1990's but needs to be revived
- **Learning representations for complex reasoning**
  - "recursive" networks [Pollack 90's] [Bottou 10] [Socher 11]
- **Integrating Deep Learning with "memory"**
  - LSTM [Hochreiter 97], MemNN [Weston 14], NTM [Graves 14]
- **Representation learning in natural language processing**
  - [Y. Bengio 01],[Collobert Weston 10], [Mnih Hinton 11] [Socher 12]
- **Better theoretical understanding of deep learning and convolutional nets**
  - e.g. Stephane Mallat's "scattering transform", work on the sparse representations from the applied math community....

Towards Practical AI: Challenges

- Applying deep learning to NLP (requires “structured prediction”)
- Video analysis/understanding (requires unsupervised learning)
- High-performance/low power embedded systems for ConvNets (FPGA/ASIC?)
- Very-large-scale deep learning (distributed optimization)
- Integrating reasoning with DL (“energy-based models”, recursive neural nets)

- Then we can have
  - Automatically-created high-performance data analytics systems
  - Vector-space embedding of everything (language, users,...)
  - Multimedia content understanding, search and indexing
  - Multilingual speech dialog systems
  - Driver-less cars
  - Autonomous maintenance robots / personal care robots

The Future: Unification
Feed-Forward & Feedback; Supervised & Unsupervised

- Marrying feed-forward convolutional nets with generative "deconvolutional nets"
  - Deconvolutional networks
    - [Zeiler-Graham-Fergus ICCV 2011]

- Feed-forward/Feedback networks allow reconstruction, multimodal prediction, restoration, etc...
  - Deep Boltzmann machines can do this, but there are scalability issues with training

Finding a single rule for supervised and unsupervised learning

- Deep Boltzmann machines can also do this, but there are scalability issues with training

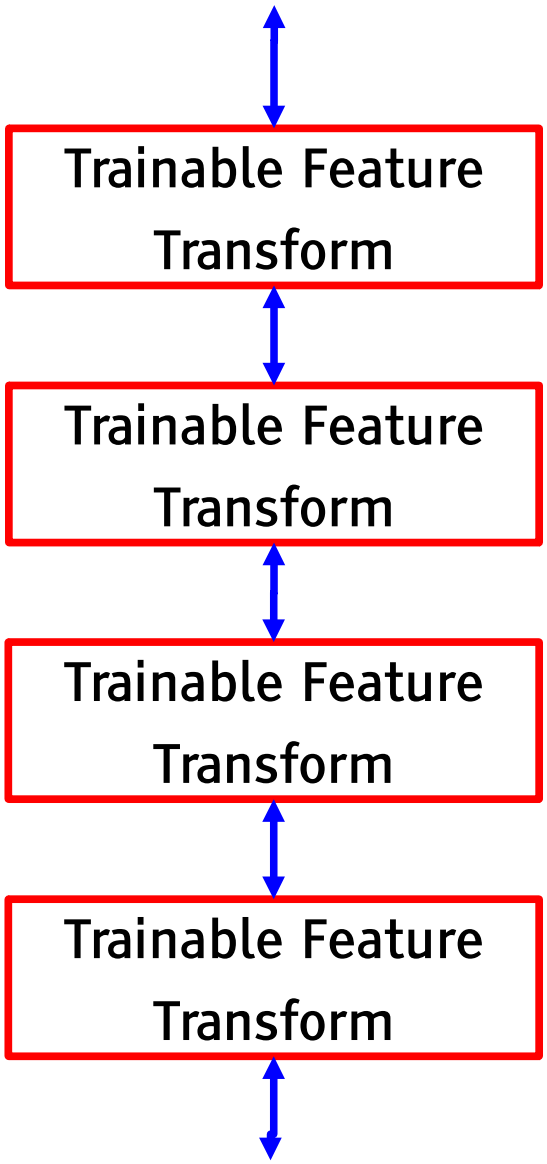

The Graph of Deep Learning ↔ Sparse Modeling ↔ Neuroscience
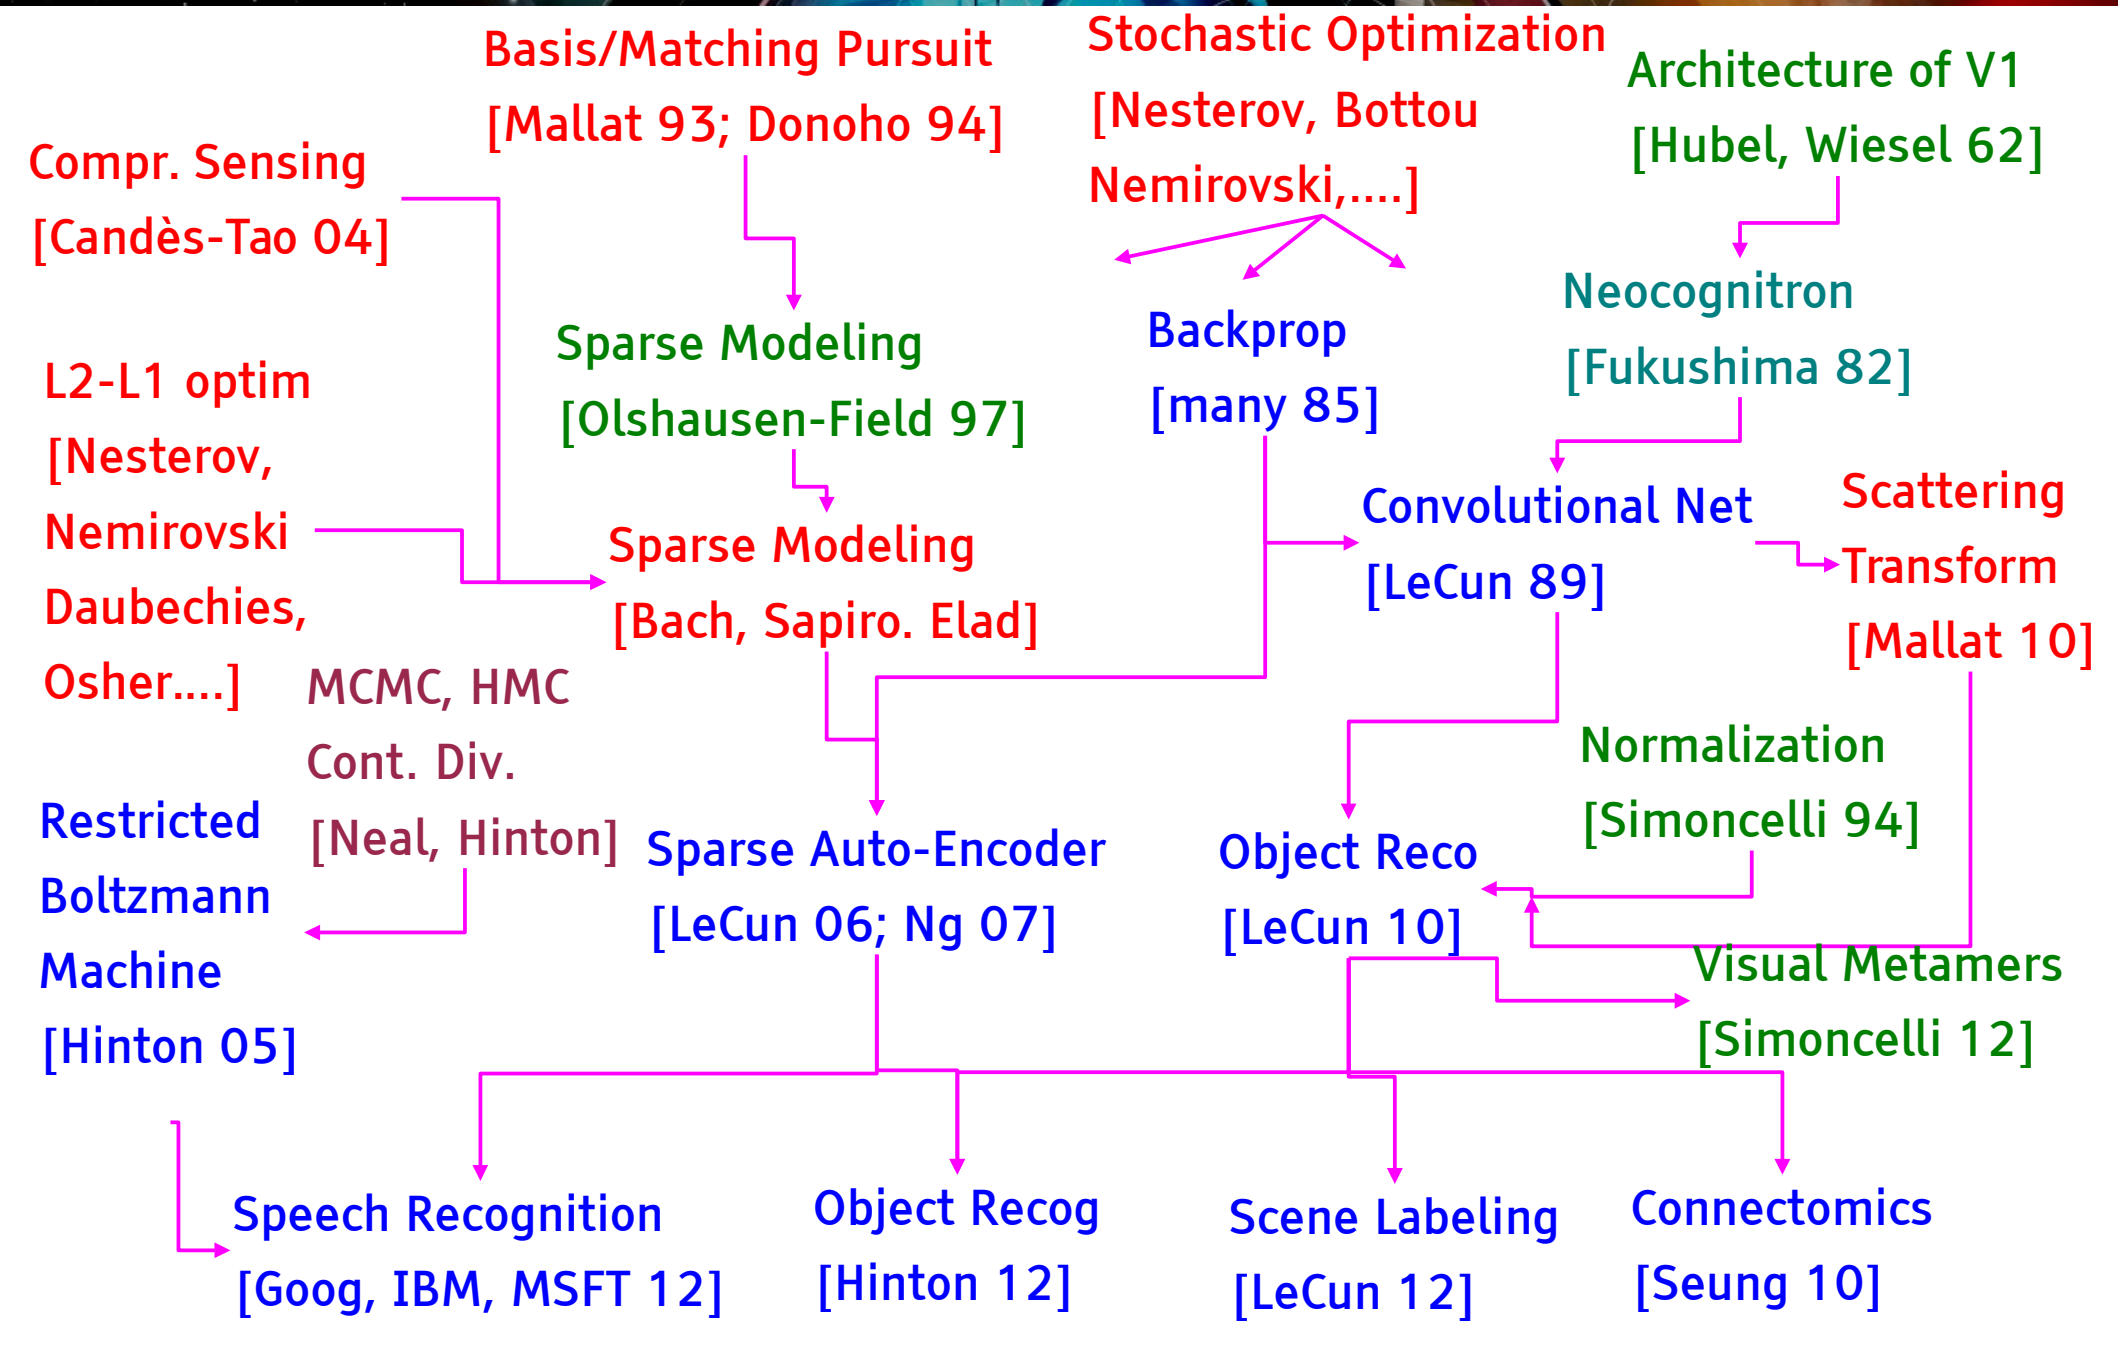
Y LeCun